Chapter 4

Magnetic Resonance Imaging

1. Introduction: Principles of Magnetic Resonance Imaging
2. Hardware
   - Magnet Systems: Current Status and Opportunities
   - Pulsed-Field MRI Systems
   - Radiofrequency Coils for MRI
   - Gradients for Dynamic MRI
3. Dynamic MR Image Reconstruction
4. Applications of Dynamic MRI
   - Blood Flow
   - Diffusion Imaging
   - Other Tissue Parameters
   - Functional Brain Imaging
   - Microscopic Imaging

This chapter highlights areas of interest related to magnetic resonance (MR) technology and its applications, particularly applications that involve dynamic magnetic resonance imaging (MRI). Among the hardware systems of an MRI device, the magnet, radio-frequency (RF), and gradient systems deserve particular R&D attention. For example, the development of suitable high-temperature superconductors (HTSs) and their integration into MRI magnets represents one potentially fruitful area of interaction between solid-state physicists, materials scientists, and the biomedical imaging research community. High-speed imaging pulse sequences, with particular focus on functional

Figure here Figure 4.1. Resolution, expressed as the reciprocal of imaging voxel volume, achieved in MR brain images by means of a volume head coil in scan times of less than 20 minutes and SNR sufficient to delineate anatomic structures.

imaging, are discussed in this chapter, as are algorithms for image reconstruction; both are promising fields of research. Significant attention is given also to two applications for which MRI has unique potential: blood flow imaging and quantification, and functional neuroimaging based on exploiting dynamic changes in the magnetic susceptibility. Finally, this chapter highlights recent technological advances toward real-time monitoring of interventional and therapeutic procedures.

MRI technology has undergone amazing strides over the last two decades, much of it due to advances from the mathematical sciences and physics. For example, Figure 4.1 demonstrates a tremendous improvement in resolution over that period. There is every reason to believe exciting progress still lies ahead.

Figure here Figure 4.2. Time domain nuclear magnetic resonance signal from volume element $dxdydz$ in an object of magnetization density $M_{xy}(t)$ in the presence of a spatial encoding gradient $\mathbf{G}$.

4.1 Principles of Magnetic Resonance Imaging

Unlike its x-ray counterparts, magnetic resonance imaging (also known as nuclear magnetic resonance (NMR) imaging) is not a transmission technique. Rather, the material imaged is itself the signal source (i.e., the macroscopic spin magnetization $\mathbf{M}$ from polarized water protons or other nuclei, such as ^{23}Na or ^{31}P). The motion of the magnetization vector of uncoupled spins, such as those for protons in water, is conveniently described in terms of the phenomenological Bloch equations:

$$\frac{d\mathbf{M}}{dt} = \gamma \mathbf{M} \times \mathbf{H} - \frac{M_{xy}}{T_2} - \frac{M_z - M_0}{T_1}\,, \tag{4.1}$$

where γ is the gyromagnetic ratio, $\mathbf{H}$ the effective field, M_0 the equilibrium magnetization, and T_1 and T_2 the relaxation times. T_1 is the characteristic relaxation time for longitudinal magnetization to align with the magnetic field: following a perturbation such as a 90° RF pulse, the longitudinal magnetization typically returns to its equilibrium value, M_0, with a time constant T_1. Likewise, T_2 is the characteristic time for decay of coherent magnetization in the transverse plane: the transverse magnetization decays exponentially with time constant T_2 to its equilibrium value, $M^{\circ}_{xy} = 0$. Both relaxation times are determined by the interaction of water or other nuclei with macromolecules in tissues. T_1 and T_2 contribute independently to the contrast between different tissues.

There is, in general, no closed-form solution to equation 4.1 (although section 14.1.6 introduces two approximate solutions). Ignoring the relaxation terms (which can be done without loss of generality when formulating the imaging equations), the steady-state solution of equation 4.1 in the presence of a static polarizing field $H_0 = H_z$ corresponds to a precession of the magnetization about the field at a rate $\omega_0 = -\gamma H_0$. Since the detected signal voltage is proportional to $d\mathbf{M}/dt$ and T_1 is on the order of 1 s, only components transverse to the polarizing field give rise to a detectable signal.

At its core, MRI exploits the field dependence of the precession frequency by superimposing a magnetic field gradient $\mathbf{G} = (\partial H_0/\partial x, \partial H_0/\partial y, \partial H_0/\partial z)$ onto the static polarizing field $H_0 = H_z$ to spatially encode information into the signal. In this manner, the resonance frequency ω becomes a function of spatial position $\mathbf{r}$, according to

$$\omega(\mathbf{r}) = -\gamma(H_0 + \mathbf{G} \cdot \mathbf{r})\,. \tag{4.2}$$

If precession due to the polarizing field (the first term in equation 4.2) is ignored, the complex MR signal is seen to evolve as $\exp(-i\gamma \mathbf{G} \cdot \mathbf{r}t)$. Following excitation by an RF pulse, and ignoring relaxation, the time-dependent signal $dS(t)$ in a volume element $dxdydz$ becomes (Fig. 4.2)

$$dS(t) = M_{xy}(\mathbf{r}) \exp[-i\gamma(G_x x + G_y y + G_z z)t] dxdydz\,. \tag{4.3}$$

It is convenient to express the signal as a function of the spatial frequency vector

$$\mathbf{k}(t) = \gamma \int_0^t \mathbf{G}(t')\, dt'\,, \tag{4.4}$$

where $\mathbf{G}(t')$ is the time-dependent spatial encoding gradient defined above, which imparts differing phase shifts to spins at different spatial locations. Equation 4.4 shows that the time variation of the spatial frequency vector is determined by the integral over the gradients; the time sequence of the RF pulses and magnetic field gradients is known as the "pulse sequence" of the MRI acquisition.

Rather than expressing the MRI signal in terms of the magnetization, which is modulated by the relaxation terms in a manner specific to the excitation scheme used, it is customary to describe the signal as a function of the density of spins in the tissue. For an object of spin density $\rho(\mathbf{r})$ the spatial frequency signal $S(\mathbf{k})$ thus is given by

$$S(\mathbf{k}) = \int_{\mathbf{r}_1}^{\mathbf{r}_2} \rho(\mathbf{r}) e^{-i\mathbf{k}(t)\cdot\mathbf{r}}\, d\mathbf{r}\,, \tag{4.5}$$

with integration running across the entire object. Pictorially, the spatial frequency may be regarded as the phase rotation per unit length of the object the magnetization experiences after being exposed to a gradient $\mathbf{G}(t')$ for some period t.

Figure here Figure 4.3. Simplified block diagram of a typical MRI system.

One further recognizes from equation 4.5 that $S(\mathbf{k})$ and $\rho(\mathbf{r})$ are Fourier transform pairs, and thus

$$\rho(\mathbf{r}) = \frac{1}{2\pi} \int_{\mathbf{k}_1}^{\mathbf{k}_2} S(\mathbf{k}) e^{i\mathbf{k}(t)\cdot\mathbf{r}}\, d\mathbf{k} \,. \tag{4.6}$$

Following an RF pulse, the transverse magnetization is subjected to the spatial encoding gradient $\mathbf{G}(t)$, which determines the spatial frequency vector $\mathbf{k}(t)$ according to equation 4.4. The manner in which k-space is scanned and the path of the k-space trajectory are determined by the waveform of the gradients $\mathbf{G}(t)$ and the sequence of RF pulses.

In contrast with the signal used in computed tomography, the MRI signal is sampled in the spatial frequency domain, rather than in object space. The most common technique involves alternate application of two or three orthogonal gradients. Signal detection and sampling in the presence of a constant, so-called readout gradient applied along one dimension then produces a rectilinear grid of data from which the image pixel amplitudes are obtained by the discrete Fourier transform. Other imaging schemes involve radial or spiral coverage of k-space.

A block diagram of a typical MRI apparatus is given in Figure 4.3. The heart of the apparatus is the magnet system, typically a superconducting solenoid. The second system is the transmit/receive assembly, typically consisting of a transmitter and power amplifier to generate the RF voltage to be fed into the transmit RF coil, thereby creating a circularly polarized RF field. The latter produces transverse magnetization, which in turn induces an RF voltage in the receive coil (which may actually serve for both transmission and reception). The ensuing signal is amplified and demodulated in the receiver electronics. Unique to the MRI device is the gradient system, which permits generation of the previously defined time-dependent gradient fields needed for spatial encoding. Both the transmit/receive and gradient systems are under the control of a data acquisition processor that ties into the main central processing unit.

4.2 Hardware

4.2.1 Magnet Systems: Current Status and Opportunities

The core of an MRI apparatus is the magnet that generates the field for nuclear polarization. The magnetic field for MRI must be extremely stable (0.1 ppm per hour) and uniform (10 to 50 ppm in a sample volume of 30- to 50-cm diameter). Magnet types in current use are of the superconducting, resistive, and permanent design. The large majority of MRI units use superconducting magnets, which provide fields of high strength and stability. They are, however, more expensive to manufacture and are cost-effective only for fields of 0.3 tesla (T) and higher. Most currently produced magnets are based on niobium-titanium (NbTi) alloys, which are remarkably reliable. These typically require a liquid helium cryogenic system with regular replenishment of the cryogens to keep the conductors at superconducting temperature (approximately 4.2 K).

In recent years the cryogenic efficiency of such magnets has been improved so that many systems now require replenishment at 9-month intervals only, compared to monthly intervals 10 years ago, and approaches eliminating the liquid helium and using conduction cooling have been proposed. Liquid nitrogen, which previously was replenished weekly, has now been completely eliminated in some systems in favor of electrically driven refrigerators. Advances in active shielding of external magnetic fields, together with creative magnetic, mechanical, and cryogenic designs, have reduced the physical volume of these magnets by typically 40% over the past decade, and the once-delicate cryostats are now rugged enough so that mobile MRI installations can travel the world's highways without incident.

Permanent and resistive magnets, both iron core and air core, are cost-effective for field strengths below approximately 0.5–0.3 T. Permanent magnets, in particular, can have very low operating costs. While permanent magnets have a weight penalty, novel designs have reduced this as well, particularly when compared to the total weight of the magnet and passive ferromagnetic shield in some other designs.

Because of its ductility and ease of fabrication, the solid-state alloy of NbTi is the usual conductor used in superconductivity applications. The only other alloy used commercially is Nb_3Sn, a brittle compound but one whose higher critical temperature allows it to be used in higher-field systems or at temperatures higher than 4 K. It has been used in several novel conduction-cooled open-architecture MRI magnets developed to test interventional and therapeutic applications of MRI (see Chapter 12). These systems operate at approximately 10 K. A natural next step in the evolution of MRI technology might be to wind MRI magnets with a high-temperature superconductor (HTS), one of a class of materials with critical temperatures above approximately 25 K. The first of these materials were

developed by Bednorz and Muller in 1986, and ones with critical temperatures as high as 164 K have been reported. Basically ceramic alloys, these materials are constructed with powder metallurgical or thin-film techniques and are inherently brittle (similar to Nb_3Sn). Considerable strides have been made in the manufacture of long production lengths of bismuth-strontium-calcium-copper oxide/silver (BSCCO/Ag), which can carry higher currents than can the low-temperature superconductors at 4.2 K. In general, HTSs appear capable of carrying more current than do low-temperature superconductors at the latter's operating temperatures in fields over 14 T.

A critical component of the superconducting magnet is the leads that are required to initially energize the magnet. If the magnets are persistent, these leads are retracted in order to reduce to a minimum the heat load on the system. If the magnets are conduction cooled, the power leads and any instrumentation connected to the coil are situated in the vacuum space with the windings. Because of the difficulty of disconnecting such leads in a vacuum, conduction-cooled coils usually use permanently connected leads, creating a significant heat load. There is promise that using HTS electrical leads could reduce such loads because of their natural low thermal conductivity.

Another critical technology is the manufacture of true zero-resistance joints for connecting individual lengths of the magnet coils and for the persistent switch used to disconnect the magnet coils once the magnet has been energized. Creation of such components for NbTi magnets is routine but is more involved for Nb_3Sn, although the use of this material in NMR spectrometers indicates that the joint technology does exist.

4.2.2 Pulsed-field MRI Systems

MRI systems require a relatively high magnetic field for polarizing the magnetic moments to provide an adequate signal for signal-to-noise ratio (SNR) considerations. This polarizing operation, of itself, is relatively noncritical. Large variations, of the order of 10%, can readily be tolerated since such a variation merely shades the resultant image and can be easily compensated for if desired. However, when this same field is used for the readout operation, as is done in all existing systems, the requirements are drastically different. The field must now have homogeneities of the order of 1.0 ppm to avoid signal loss and distortion. The signal loss stems from variations in frequency within an imaging voxel[1] where the resultant destructive phase interference causes enhanced signal decay, significantly reducing the sampling time available. These problems are particularly severe in high-speed systems that attempt to cover large amounts of k-space following each excitation. The limited readout time restricts the amount of coverage and the SNR, making real-time performance impractical.

In pulsed-field systems the functions of polarization and readout are separated, enabling a number of significant performance and cost advantages. A pulsed field is applied and left on for a time comparable to the relaxation time T_1, to polarize the magnetic spins. Since the magnet has considerable thermal inertia and is on only during the imaging interval, the peak field can be made relatively high, of the order of 1.0 T, without excessive temperature rise. Thus the very use of a pulsed field has a significant SNR advantage over static resistive magnets. In addition, the time of the pulse can be used for providing T_1 contrast by determining the degree of polarization buildup for the different materials. Of course T_2 contrast is established in the usual way by varying the time over which the origin of k-space is scanned during the readout. As indicated, this pulsing operation is highly noncritical as to homogeneity. It therefore enables the use of low-cost simple magnets. The magnets now being used are either very large solenoids or vertical field magnets with pole pieces much larger than the imaged region. These have high stored energy and result in very limited accessibility to the patient for guided biopsies and/or therapeutic procedures. Pulsed magnets enable the use of simple configurations such as a pair of coils or a "C"-shaped core having a cross section comparable to the region of interest. This simplicity enables the construction of specialized machines for joints, head, breast, heart, and so on. The readout bias field following the pulse is made very low so that its variations will result in negligible T_2^* or distortion considerations.[2] Susceptibility-induced variations (which scale with readout field strength) should be negligible.

Although the readout bias field is relatively noncritical, it can be a major consideration in a simple low-cost design. For example, if the readout field were 1,000 times smaller than the polarizing field, for comparable performance its homogeneity should be 1 part per 1,000 as compared to 1 ppm. Although this ratio is readily realized in a large solenoid, it remains quite difficult to achieve in a simple configuration such as a pair of coils positioned on opposite sides of the body. To achieve essentially complete immunity to inhomogeneity of the readout bias field, comparable to the immunity of the pulsed polarizing field, an oscillating bias field is used.

Perhaps the biggest problem created by this approach is the radio-frequency receiving system. The desired result is to obtain an SNR commensurate with the size of the polarizing field. This will happen only if the noise generated is due primarily to body losses, where other losses, including those of the coil, are negligible. The problem is that,

[1] **V**olume **el**ement, the higher-dimensional analog of the more-familiar pixel (**pi**cture **el**ement) that is the basic unit of a planar image.

[2] T_2^* is the effective transverse relaxation time, which determines the rate of decay of the signal in the presence of magnetic field inhomogeneity.

at very low readout frequencies, the body losses are significantly reduced, thus making coil noise relatively more important. As a result systems using very low readout frequencies require low-loss coils. These can include coils made of cooled materials and high-temperature superconductive coils. The use of these low-loss, high-quality-factor coils introduces another problem, that of insufficient bandwidth.

The small signal also results in an increased shielding problem since external interfering signals are more significant. One solution to this problem is the use of gradiometers, coil configurations that respond solely to field variations and ignore uniform fields arising from more remote sources.

4.2.3 Radio-frequency Coils for MRI

Radio-frequency coils constitute the essential means of stimulating the spin systems and are the doorway for the sensing of MR data. Several research opportunities are presented by the need to improve the quality of the data and the rate of data acquisition. The data quality could be improved with better coil designs and, in some applications, cooled or superconducting coils, which result in lower noise. Better data quality could also be facilitated with multiple-acquisition coils operating in parallel.

Computational Design of RF Coils Historically, RF coils have been constructed according to basic design principles learned through experience by the individual investigator. More sophisticated computational RF coil design is an area for future growth, for several reasons:

- MRI is an intermediate field problem, and such fields are difficult to simulate even for quasi-static analyses.
- Very high frequencies are beyond the quasi-static regime of Maxwell's equations. Wavelength effects inside the patient are important and must be considered. Designing transmitters and receivers based on an understanding of dielectric and wavelength effects, so as to enable uniform excitation and detection, is beyond the scope of quasi-static arguments.
- Accurate estimates of power deposition in patients are needed for high-field, high-power imaging sequences. Since neither experimental measurements nor first-order calculations can accurately predict local power deposition, accurate field calculations are required as a basis for developing safe operating guidelines.
- Very high frequency coils are still tuned and matched by trial and error. As in the design of lower-frequency resonators, it would be extremely useful to have a set of independent controls for the coil's different oscillatory modes that produce circularly-polarized fields.

All of these applications require the availability of a true three-dimensional solution to the time-varying fields. As yet, there has been little published on MRI applications of finite-element calculations. A few abstracts have described some small-scale problems, but none have attacked a full size problem. Generally this work has not taken a full three-dimensional approach; usually one dimension in the problem has been very coarsely sampled in the analysis.

Assuming that local heating effects in a body need to be quantified, a finite-element calculation with elements on the order of 1 cm^3 would be needed. A head-size coil in a shield might require something on the order of 100,000 to 200,000 elements. The scale of this problem is appropriate for even a modest modern workstation; with an iterative finite-element solver, something on the order of a week of CPU time might be required. This scale is just right for a research group to begin investigating. National supercomputer facilities exist that could greatly reduce this time.

Beyond the physics of solving electromagnetic problems, issues specific to MRI complicate the interpretation of results. The best methods for visualizing three-dimensional field values are not obvious. Development of such techniques will require computational tools specific to MRI, for example, means of performing a simulated MR pulse sequence and visualizing the images created by it.

Cooled Receiver Coils for MR Imaging Most SNR calculations in MRI rest on the basic assumption that the patient contributes the dominant part of the noise. This assumption is known to be incorrect in two important cases:

- When coil loading increases rapidly with increasing frequency. Consequently in low-field MRI, patient contributions to noise may not dominate.
- When coil loading decreases rapidly with the length scale of the system. As the receiver coil is made smaller, the contribution from patient loading can be made smaller than coil contributions at any field strength. For microscopy applications the coil is the dominant source of resistance.

High-temperature superconducting materials that contribute little noise when operated at liquid nitrogen temperatures deserve to be evaluated for use in MRI receiver coils. The most promising technology involves thin film HTSs. These materials can be used in large static fields provided that the plane of the thin film is parallel to the field. The limitations of thin film technology are those imposed by the vapor deposition process used to create them. Very large coils are at present not possible; the practical upper limit is about 8 cm. Alternative possibilities are tape HTSs or the use of conventional copper coils at liquid nitrogen temperature.

Use of Multiple Receivers While the use of multiple acquisition coils operating in parallel could improve data quality, the technological requirements for a practical multi-coil array are formidable: beyond the issues of building a decoupled array of surface coils, there are considerable demands on the system hardware to be able to deal with data rates that are an order of magnitude higher than those from a conventional receiver coil. A flexible and powerful system architecture is needed that allows for the acquisition and manipulation of larger data streams. An active research area is the multiplexing of several RF channels through one wideband receiver, digitizer, and computer. Multiplexing is possible in both the frequency and time domains. Noise crosstalk between channels must be prevented by tight filtering near the RF coils. Crystal filters and filters built with HTS high current resonators are contenders for this task.

The simplest approach for the combination of data from the array is based on a linear combination of the individual images into a composite image. Provided the receiver coil sensitivities are known and a collection of equal-noise images exists, then the least-noise combined image is a sum of images weighted by their sensitivities. However, achieving reliable and noise-free estimators of individual coil sensitivities remains a challenge that calls for more sophisticated techniques, possibly statistical ones.

Alternatively, there are several adaptive techniques for image combination that act strictly on the original images without any need for filtering. Principal component analysis (PCA) can take a collection of original images with different noise levels and synthesize a composite image that is the noise-optimum linear combination. A collection of surface coil images requires a linear combination in which the linear weighting is spatially dependent. To calculate something like a PCA combination, one would require a means of performing an eigenvector analysis on a small segment of the entire image, repeating this process across the image, and then interpolating between the different subimages. Using the interpolation scheme is essentially the same as working from estimators based on filtered images, but the advantage is that an assumption of equal noise levels is no longer necessary. One could then produce an optimal linear combination of surface coil images derived from input images with greatly different noise characteristics.

Multiple coil technology might also be used to speed up the MR imaging sequence. One can argue from simple heuristic principles that acquisition of signal from two (or more) receiver coils yields twice as much information per phase encoded signal and that we ought to be able to use the additional information to reconstruct an image based on half the number of phase encoded signals. One method of implementing this uses receiver coils constructed with a known sensitivity. It requires that one coil be a standard uniform receiver of the type common in MRI; the second is an RF gradient coil. Fortunately these coil sensitivities are easy to implement in a conventional horizontal-field superconducting magnet: one coil is the standard Helmholtz pair, and the second is a Helmholtz pair with opposed windings. An imaging acquisition consists of a standard phase encoded sequence, but every other acquisition is skipped. The skipped phase encoded echoes can be calculated from the acquired data on nearby steps in Fourier space from each of the separately received data sets. Possibilities for extending these ideas to multiple receiver coils provide interesting research challenges.

4.2.4 Magnetic Field Gradients

Spatial encoding in MRI typically uses spatial variations in the static or radio-frequency magnetic field. Because such variations are small compared to those in the static magnetic field B_0 used for MRI, a set of "gradient coils" can be used to produce a very nearly linear gradient along any direction. Gradients in the radio-frequency magnetic field, although useful for some purposes, are not discussed here.

Besides being the essential element for spatial encoding, the gradient system of the MRI scanner is responsible for encoding dynamic information related to flow, diffusion, or modulation of magnetization for spatial tagging. In many applications, the quality of the data obtained is currently limited by hardware-imposed limitations on the gradient amplitude and/or rate of change (slew rate). As an example, such limitations hamper echo-planar imaging (EPI), where ideally all the k-space data would be acquired following a single excitation over a period on the order of T_2. Limited gradient amplitude and switching rates currently restrict the temporal and spatial resolution of EPI.[3]

[3] For background on EPI see, for example, the paper by Cohen and Weisskoff listed in the suggested reading for this section.

A second example is diffusion imaging, which is limited primarily by gradient amplitude to very long echo times on conventional systems. When high-performance gradient hardware is available, some applications are significantly limited by restrictions on maximum permissible field slew rates in humans ($dB/dt < 20$ T/s). Such restrictions are in place to prevent nerve stimulation by induced eddy currents.

Local versus Whole-Body Gradients The MRI scanner's gradient system consists of two major components, the gradient coil, which is a set of windings that determines the spatial distribution of current, and the gradient power supply, which typically is a digitally controlled audio amplifier that supplies power to the coil. Stronger and faster switching gradients can be achieved by improving either of these components. Local gradient coils (LGCs), which are smaller than the whole-body gradient coils that are built into clinical imaging systems, are designed for specific purposes or areas of anatomy. Their smaller size gives them three advantages over whole-body gradients. The first is higher performance in the form of stronger gradient amplitudes per unit current, and higher gradient slew rates per unit voltage. The second is that because of their smaller size, LGCs can create higher gradient slew rates for a given imposed limit on field slew rate—a more fundamental advantage in that LGCs, in the absence of power limitations, can achieve stronger gradients with faster switching than can whole-body gradient systems. The third advantage is lower cost, now approximately an order of magnitude lower.

High-powered whole-body gradient systems generally fall into one of three categories: high-powered linear amplifiers, switching systems, and resonant systems. High-powered linear systems are self-explanatory. Both switching systems and resonant systems use energy storage techniques to minimize power dissipation. Switching systems generally use semiconductor switches to step between current states very rapidly, while resonant systems use an L-C circuit with the gradient coil as the inductor to efficiently generate sinusoidal current waveforms. A few of these approaches have been used successfully to design and implement fast imaging systems capable of, for example, single-shot echo-planing imaging. The primary advantage of this approach is that the hardware is not specific to the anatomic region of interest, allowing imaging of various regions without requiring a change of gradient hardware.

The remainder of this section focuses on LGCs.

Design Considerations Many factors must be taken into account in the design of LGCs, including coil geometry, coil size, anatomy of interest, desired gradient strength, efficiency (gradient per unit current), inductance, desired slew rate, eddy currents, gradient uniformity, forces and torques, heat dissipation, and nerve stimulation.

Geometry: The geometry of LGCs, like that of local RF coils, can be divided into three categories: whole-volume coils, which completely surround the anatomic structure of interest; partial volume coils, which partially surround the structure; and surface gradient coils, which are placed against one side of the structure.

The advantages gained by the use of an LGC depend strongly on the size of the coil. For a coil of a given design, with a fixed number of turns, the efficiency scales inversely with the size squared while the inductance scales with the size, indicating that the size of the coil should be as small as possible. The lower limit on the size of the coil is set by three other factors: the coil must have a region of usable gradient field uniformity that encompasses the anatomic region of interest; the coil itself must be large enough to accommodate the anatomic region of interest; and sufficient room must be left for the RF structure and return flux to allow dominant loading of the RF coil by the subject. The largest of these three lower limits determines the optimum size of the coil.

Eddy Currents: Eddy currents are currents that are induced in conducting structures due to time-varying magnetic fields. Gradient eddy currents effectively limit the field slew rate by creating transient fields that usually oppose the desired gradient fields. Structures in an MRI scanner that may support eddy currents include the cryostat, the RF coil, and the RF shield. For whole-body gradient coils, eddy currents generated in the cryostat are usually of the greatest concern because the cryostat is physically close to the gradient coil and has very low resistance. In general, LGCs are much smaller and therefore couple much less strongly to the cryostat. However, RF coils or shields that have low-resistance closed-DC current paths in close proximity to the gradient coil can support eddy currents and degrade the gradient field waveforms. This problem must be considered carefully in the design of RF coils and shields for use with LGCs.

Linearity: Because the gradient fields directly map the spatial coordinates of physical space to the coordinates of the image, perfect gradient field uniformity is desired. However, if the gradient fields are not perfectly linear but are known, spatial transformations can be applied to remove warp from the images. This process can partially alleviate the problems associated with using nonlinear gradients, but several problems remain. In two-dimensional imaging, the most common type, if the slice-select gradient is nonlinear, then information is obtained from tissue outside the desired slice and the slice cannot be correctly reconstructed. If three-dimensional data are collected, or if the slices are contiguous, then three-dimensional spatial transformations can be used. These transformations are straightforward but increase image reconstruction time. Because nonlinearities in the gradients cause variations in

the voxel sizes, the signal intensity, SNR, and spatial resolution will be modulated by these variations, potentially adding to the difficulty of image interpretation.

Lorentz Forces: The currents carried by a gradient coil interact with the static field, causing the coil to experience large forces. In a uniform static field, a closed current loop cannot experience net forces, but it can experience net torques. If these torques are not internally balanced, the coil must be mounted in a very strong and rigid mechanical structure that is tied to the structure of the main magnet in order to prevent large-amplitude vibrations of the coil. Torques on an unbalanced head coil in a 1.5-T field are on the order of 100 N-m. Internal balancing of the torque in a coil can be achieved by symmetry or by explicit nulling of the torque moment in an asymmetrical design.

Design Techniques: Numerous mathematical techniques have been used to design gradient coils, many of which are described in a 1993 review article by Turner listed in the recommended reading below. Most designs to date describe current patterns on the surface of one or two cylinders, as this is the most natural geometry for whole-body patient access, and for LGCs a natural geometry for the head and extremities. Early designs consisted of simple pairs of loops for longitudinal gradients and saddle coils for transverse gradients. Now, a more elegant approach has been introduced that explicitly minimizes inductance while creating the desired fields within the region of interest. Inductance is a quantity that is difficult to minimize using other techniques because it is a nonlinear function of the current density. For a given coil geometry, an appropriate transform applied to the current density linearizes the inductance, greatly simplifying the problem of inductance minimization. For planar and cylindrical geometries, simple two-dimensional Fourier transformation serves this purpose. In the transformed space the inductance is a linear function of the transform components, and a Lagrange multiplier technique can be used to satisfy the desired field constraints while minimizing inductance. After deriving a current density pattern in transform space, a reverse transform is applied to produce a current density pattern in physical space, and approximation by discrete current paths gives a final design. Several variations of the minimum inductance technique have been introduced to apply the technique to different geometries and to use the same principles to minimize power dissipation instead of inductance.

Another class of design techniques uses numerical methods to optimize combinations of gradient uniformity, inductance, and coil efficiency. Using either a parameterized or a point-by-point description of the current path, mathematical techniques such as gradient descent, simulated annealing, Monte Carlo, and simplex can be used to minimize a chosen cost function. Numerical methods are in general more computationally intensive than analytical techniques, but they allow for incorporation of arbitrary constraints in a simple manner, and they can directly produce final designs, avoiding errors that may be introduced by the approximation of continuous current densities by discrete wires.

Applications Besides being an integral part of the spatial encoding process, high-performance gradient systems have many other applications. In angiography, stronger and faster gradients allow for shorter flow-compensating gradients and hence shorter echo times, T_E. These conditions decrease flow-related dephasing and increase the detected signal. In saturation tissue tagging, such as myocardial tagging to observe cardiac wall motion, stronger gradients can improve tag profiles and decrease tagging time, thus generating sharper tagged grids in the final images. For imaging of very short echo time T_2^* regions, it is necessary to employ strong gradients in order to ensure that the imaging gradients dominate the static field inhomogeneities. Short T_2^* occurs primarily in areas with strong susceptibility gradients such as in lung tissue.

Bioeffects For MRI in humans, the fundamental limitation on the performance of gradient systems is the need to avoid nerve stimulation. For the application of single-shot echo-planar imaging (EPI) of the head, for example, a three-axis gradient coil enables the use of EPI in an otherwise conventional commercial scanner. Using standard gradient amplifiers, the 20-T/s limit imposed by the U.S. Food and Drug Administration can be exceeded with this coil, so that EPI is no longer hardware limited. A part of the coil optimization question that has not been addressed specifically is how to design a coil with appropriate field distribution, access, efficiency, and so on that also explicitly minimizes the peak field to which the subject is exposed. This is an added dimension to the coil design problem that will have a different solution for every imaging region of interest and every anatomical site.

4.2.5 Research Opportunities for MRI Hardware

Magnet Systems

- Development of *economical* higher-temperature superconducting magnets, using Nb_3Sn and other higher-temperature materials.
- Development of designs for *economical* magnets for special applications, including specific anatomic parts, or for special disciplines such as therapy.

- Development of designs for more economical magnetic field shielding that would allow retrofitting of magnets into existing diagnostic, interventional, and surgical rooms.

Pulsed-field MRI

- Development and validation of strategies for signal readout that minimizes introduction of additional noise and interference.
- Development of means for energy recovery during collapse of the polarizing field.

RF Coils

- Design of uniform transmitters and receivers that include dielectric and wavelength effects.
- Development of methods for three-dimensional modeling of RF fields.
- Design and evaluation of high-temperature superconducting coils and associated cooled pre-amplifiers for low-field imaging.
- Design of SNR-efficient high-speed image combination and reconstruction techniques for multicoil arrays.

Gradient Systems

- Design of very short head gradient coils using current return paths at a greater diameter than that of the primary windings.
- Design of acoustically quiet gradient coils.
- Design of head gradient coils with good subject access for visual and auditory task presentation.
- Optimization of head immobilization devices compatible with head gradient coils.
- Development of methods for electrical decoupling of RF coils and gradient coils in close proximity.

4.2.6 Suggested Reading Related to MRI Hardware

Magnet Systems

1. Stekly, Z.J.J., and Gregory, E., in *High-Temperature Superconducting Materials Science and Engineering*, Pergamon Press, Tarrytown, N.Y., 1994.
2. Stekly, Z.J.J., and Gregory, E., in *Intermetallic Compounds: Principles and Practice*, J. Westbrook and R.L. Fleisher, eds., John Wiley & Sons, New York, 1994.

Pulsed-field MRI

3. Macovski, A., and Connoly, S., Novel approaches to low-cost MRI, *Magn. Reson. Med.* **30** (1993), 221–230.
4. Mansfield, P., and Morris, P., *NMR Imaging in Biomedicine*, Academic Press, New York, 1982.
5. Morris, P.G., *Nuclear Magnetic Resonance Imaging in Medicine and Biology*, Oxford University Press, New York, 1986.

RF Coils

6. Carlson, J.W., and Minemura, T., Imaging time reduction through multiple coil data acquisition and image reconstruction, *Magn. Reson. Med.* **29** (1993), 681–688.
7. Ra, J.B., and Rim, C.Y., Fast imaging using subencoding data sets from multiple detectors, *Magn. Reson. Med.* **30** (1993), 142–145.
8. Roemer, P.B., Edelstein, W.A., Hayes, C.E., Souza, S.P., and Mueller, O.M., The NMR phased array, *Magn. Reson. Med.* **16** (1990), 192–225.
9. van Heteren, J.G., James, T.W., and Bourne, L.C., Thin film high temperature superconducting RF coils for low field MRI, *Magn. Reson. Med.* **32** (1994), 396–400.

Gradient Systems

10. Cohen, M.S., and Weisskoff, R.M., Ultra-fast imaging, *Magn. Reson. Imaging* **9** (1991), 1–37.

11. Kwong, K.K., Belliveau, J.W., Chesler, D.A., et al., Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation, *Proc. Natl. Acad. Sci. USA* **89** (1992), 5675–5679.

12. Ogawa, S., Tank, D.W., Menon, R., et al., Intrinsic signal changes accompanying sensory stimulation: Functional brain mapping with magnetic resonance imaging, *Proc. Natl. Acad. Sci. USA* **89** (1992), 5951–5955.

13. Rosen, B.R., Belliveau, J.W., and Chien, D., Perfusion imaging by nuclear magnetic resonance, *Magn. Reson. Quart.* **5** (1989), 263–281.

14. Rzedzian, R.R., and Pykett, I.L., Instant images of the human heart using a new, whole-body MR imaging system, *Amer. J. Roentgenol.* **149** (1987), 245–250.

15. Turner, R., Gradient coil design: A review of methods, *Magn. Reson. Imaging* **11** (1993), 903–920.

16. Turner, R., Minimum inductance coils, *J. Phys. E: Scientif. Instrum.* **21** (1988), 948–952.

17. Turner, R., Optical reconstruction of NMR images, *J. Phys. E: Scientif. Instrum.* **18** (1985), 875–878.

18. Turner, R., Jezzard, P., Wen, H., et al., Functional mapping of the human visual cortex at 4 and 1.5 tesla using deoxygenation contrast EPI, *Magn. Reson. Med.* **29** (1993), 277–279.

4.3 Dynamic MR Image Reconstruction

Dynamic magnetic resonance imaging today can be interpreted to mean either very fast or essentially continuous scanning. In practice, dynamic means fast enough to observe an object before its properties change: fast enough to image the heart before it moves or fast enough to image a tissue that is taking up a contrast agent before its relaxation properties change. As such, dynamic imaging requires the fastest imaging schemes possible. So, how does image reconstruction enter the picture? From the simplest perspective, the fastest possible reconstruction scheme is desired, which today is probably the fast Fourier transform (FFT). However, when data are collected quickly in an MRI experiment, the signal-to-noise ratio (SNR) is reduced. So a reconstruction method with the highest possible SNR is desired as well. Another factor is the resolution. Physical characteristics will also affect the choice of resolution. Finally, and most importantly, how do all these needs affect the way the MRI data themselves are collected so as to optimize the image quality and information content? Can the concepts of resolution, SNR, and speed be uncoupled? Clearly, this image reconstruction problem represents a challenge involving both the mathematics and the physics of MRI.

In this brief outline of research opportunities related to real-time MR image reconstruction, several approaches are considered: partial Fourier reconstruction, parametric estimation, Bayesian analysis, wavelets, image filtering, echo-planar imaging, and the general question of k-space coverage. One particularly important point to make here is that most of these methods have focused on one-dimensional problems with some brief forays into higher dimensions. In MRI these higher dimensions provide a wealth of data that require alternative methods of reconstruction.

4.3.1 Partial Fourier Reconstruction

In the simplest version of MR image reconstruction, data are collected in a complex form symmetrically about the origin. In k-space notation, the data run from $-k_{\max}$ to $k_{\max}$ in discrete steps. For example, when the FFT is used, $2n$ points are collected for $-n \le i \le n-1$ with $k(i) = i\Delta k$. To save time, data can be collected in principle for just $0 \le i \le n-1$ and, when the object to be reconstructed is real, the data can be complex-conjugated and a conventional complex FFT done to recover the real object. In practice, the reconstructed image data are complex-valued due to factors such as main field inhomogeneity and flow effects. Discussed below are four specific strategies that attempt to deal with this more general situation. These discussions ignore Gibbs ringing, but it is included in section 4.3.2.

Predominantly One-sided Data Collection The most common way used to save time is to collect just m points before the echo and n points after the echo. The time to acquire the data is thus reduced to a fraction $(m+n)/2n$. The loss in SNR is the square root of this fraction. Unfortunately, when the object is complex, the operation of complex conjugation is foiled. One approach used was to assume that the object had only low-spatial-frequency phase variations that were well defined within the data points $-m \le i \le m$. A low-spatial-frequency image was then reconstructed and the phase from this image used to rotate the additional information into the real part by

the above convolution. This common procedure is technically correct only for small phase angles. In practice, a projection onto convex sets (POCS) approach can be used to obtain a more stable solution to $\rho(x)$.

The procedure can actually be restated as a matrix problem and a solution to the inverse can be found. The disadvantage to this approach is that matrix inversion is time-consuming. The advantage is that the ill-posedness of the problem can be better understood. To date, though, the general problem of correcting for a high-spatial-frequency phase error has not been fully explored. The reconstruction errors that manifest themselves in the presence of these phase errors show up as an oscillatory signal response about the region where the error has not been accounted for. These reconstruction errors can be quite large and often cause bright or dark spots in the MR image. They are thus very detrimental to the use of one-sided data collection, making quite problematic its application to data obtained with techniques that tend to produce phase errors, such as gradient echo sequences or sequences that are not flow compensated.

Predominantly Every Other Point There are indeed many other ways to attempt to optimize the collection of the data to deal with the above problem, but one that is of particular interest is that which collects the same data in the center (i.e., from $-m \leq i \leq m$) and then from all odd points for $-n \leq i < -m$ and all even points for $m < 1 \leq n-1$. Collecting the data in this way has the advantage that high-spatial-frequency information is collected throughout what would have been the conventional data set and hence can accommodate field inhomogeneity errors that are either positive or negative. The blurring artifacts vanish completely and are replaced instead with aliasing, or half-frame displaced ghosts. The images are sharp, but any remnant phase error not accounted for leads to an interpolation error that manifests itself as ghosting rather than blurring.

There is therefore often more than one way to attempt to collect the same data, and the method of choice may be determined by some physical principles related to how the data are collected or to the type of artifacts. For example, for the shortest possible echoes, the asymmetric data from predominantly one-sided data collection would be the best choice, whereas if field inhomogeneities cause unacceptable artifacts, the method described in the preceding paragraph is the preferred one.

Collecting Multiple Echoes Adding a little more physics to the problem, one can ask the question, How can I supplement the lack of information available from, say, a very short asymmetric echo? One possibility is to add a second echo after this that may or may not itself be symmetric. The second echo will have worse phase errors than the first but the phases of the two are simply related. However, one must then deal with another reason why it is so difficult to carry out partial Fourier reconstructions: the data $s(k)$ are the Fourier transform of an object that is spatially varying, hence, all points in $s(k)$ contain information about all points in $\rho(x)$; deconvolving this information is computer intensive. It is possible to use this second echo along with a POCS-like iterative approach to more accurately solve this problem.

Two- and Three-Dimensional Extensions A largely ignored problem is the application of these data extension methods to two and three dimensions. A tremendous increase in time savings is possible in three-dimensional imaging if sampling is reduced in both phase encoding and slice encoding directions. The key issue is whether there is sufficient information available in any data acquisition scheme to uniquely solve the problem.

4.3.2 Reduced Gibbs Ringing

In the presence of a discontinuity, the Fourier transform yields an estimate that has a "ringing" at the discontinuity. This is true for any number of sampled points, and the Gibbs ringing is essentially scale invariant. This means that the ringing is always present at the same pixel values with similar amplitude. However, as the number of points in an MR image of given size increases, so also does the resolution, and so the points at which the ringing occurs are squeezed spatially closer to the discontinuity and thus become less likely to contribute to misinterpretation of the image. Recall that this famous Gibbs ringing or truncation artifact decays very slowly (as $1/x$ in the image domain), and so it is considered a particularly nasty artifact, one that can obscure information or mimic structure that is not really present physically. Therefore, quite some effort has been made in the MRI signal processing world to eliminate or reduce these effects.

Iterative Sigma Filtering One approach to reducing Gibbs ringing is to apply a nonlinear filter to determine the major edges (or discontinuities) in an image and then apply a sigma filter (or some other similar nonlinear filter) to smooth the image outside the edge. Interpolating the image by a factor of two, for example, and reconstructing a new data set with $4n$ points, leads to an extrapolated k-space data set that supposedly has recovered some of the missing high-k-space data points. This process can be repeated iteratively by modifying the filter values as a function of iteration number.

Constraint-based Methods Another approach to filtering images can be seen with MR angiographic images. Here the starting assumption is that the image can be converted to binary format: pixels can be segregated into two classes, those that depict parts of blood vessels and all others. Then resolution can be enhanced by taking partial volume effects into account to relabel some pixels. This method of filtering has been tackled rather nicely using a Bayesian approach, and more can perhaps be done with multiple image components.

Parametric Estimation Parametric estimation techniques are the most powerful methods for extracting information from an image when a priori information is available. Various approaches use linear prediction techniques and autoregressive moving average (ARMA) methods, localized polynomial approximation (LPA), and the generalized series method. These methods all attempt to recover parametric variables that describe a model. In the simplest one-dimensional example, modeling the data as a box, one would find the location of the center of the box, its amplitude, and its width. Unlike FFT methods or image processing methods, these methods have no inherent pixel size limitations. These are "super-resolution methods" that can be used to speed up image acquisition. For a box only two data points are needed!

What are the limitations of these methods and why are they not more commonly used? First, the resolution is now dependent on SNR. Second, they are often ill-posed (i.e., the condition number is large). Third, they are much slower to implement on a computer (often hundreds of times slower than the FFT). Nevertheless, with the dramatic improvement in computation speeds available today, the reconstruction times have dropped to within a few seconds in some cases, making these methods more viable in practice.

4.3.3 High-speed K-space Coverage Techniques

Moving from one-dimensional problems, one method in MRI that acquires data in a two-dimensional format and in a single shot is echo planar imaging (EPI). In fact one can cover k-space in a plethora of ways with rectilinear coverage, or in sinusoidal fashion. Each method has its own artifacts; sinusoidal coverage gives ghosts, while spiral scanning gives blurring. While speeding up imaging by simply taking fewer data points can always be done, but at the expense of resolution, these various methods are important because they attempt to overcome the resolution limits. In keyhole imaging, for example, a limited number of data are recollected as, say, a subject performs a functional imaging task where only a small portion of the image changes. Concepts such as use of shared echoes between scans to speed up spin echo or cardiac cine-imaging are similar to those described in section 4.3.1, but different enough to warrant rethinking certain data acquisition schemes that might otherwise not qualify as real-time imaging.

A final comment in this section relates to the general reconstruction question in n dimensions. First, what is a sufficient number of points to collect for a given reconstruction method? For example, for the FFT a high enough density of points can guarantee no aliasing even with non-uniform sampling. Second, how can we know how k-space is covered and, once we do, how can we correct for any deviations from the sampling we expected that are due to eddy currents or other phenomena? The first part of this question may be addressed using some features of the reconstruction method itself (such as varying the parameters until ghosting vanishes), but it is more likely answered by first collecting data that can serve as a calibration scan to determine background phase or, better yet, the exact location of all points to be sampled in k-space.

The reader should be aware that MRI is an incredibly flexible imaging modality. More often than not, thanks to the linearity of the Fourier transform and therefore of the whole data acquisition approach, many of the types of methods outlined here can be combined! By itself any one method may be unexciting; however, if each method were to give just a $\sqrt{2}$ improvement in SNR or speed, then several combined together could yield significant improvements overall.

4.3.4 Research Opportunities in Dynamic MR Image Reconstruction

Although many areas of image reconstruction have been examined as possible improvements to the FFT, few have succeeded in becoming practical enough for application in MRI. However, given ever-increasing computation speeds, the following topics may play a key role in the future of real-time scanning:

- Maximum entropy methods, although research to date has not shown them to be useful in MRI;
- Wavelets, for modifying both how data are collected and how reconstruction is performed;
- Bayesian analysis to improve signal-to-noise ratios; and
- New methods of iteration based on the search-everywhere concept of neural networks.

In all cases it is important to use the FFT as a benchmark and to ensure that the methods are stable and representative of reality under all reasonable imaging conditions for the modality.

4.3.5 Suggested Reading Related to Dynamic MR Image Reconstruction

1. Constable, R.T., and Henkelman, R.M., Data extrapolation for truncation artifact removal, *Magn. Reson. Med.* **17** (1991), 108–118.

2. Haacke, E.M., Liang, Z.-P., and Izen, S.H., Constrained reconstruction, a super-resolution, optimal SNR alternative to the FFT, *Med. Phys.* **16** (1989), 388–397.

3. Liang, Z.-P., Boada, F.E., Constable, R.T., et al., Constrained reconstruction methods in MR imaging, *Reviews of Magnetic Resonance in Med.* **4** (1992), 67–185.

4. Liang, Z.-P., and Lauterbur, P.C., An efficient method for dynamic magnetic resonance imaging, *IEEE Trans. Med. Imaging* **13** (1994), 677–686.

5. van Vaals, J., Brummer, M.E., Dixon, W.T., et al., "Keyhole" method for accelerating imaging of contrast agent uptake, *J. Magn. Reson. Imaging* **3** (1993), 671–675.

6. Wu, Z., Chung, H., and Wehrli, F.W., A Bayesian approach to subvoxel tissue classification in NMR microscopic images of trabecular bone, *Magn. Reson. Med.* **31** (1994), 302–308.

4.4 Applications of Dynamic MRI

4.4.1 Blood Flow

MRI has unique potential for imaging and quantifying blood flow. Flowing spins modulate the MR signal in two different ways. First, displacement of the bolus between successive excitations introduces some fully relaxed spins and thus an enhancement of the signal from the flowing blood relative to that of the stationary tissue. This time-of-flight effect is the basis of a major class of angiographic imaging techniques. The second effect has its origin in phase shifts imparted by the imaging gradients or specially administered flow-encoding gradients. This approach, summarized under the term "phase-contrast" (PC) MRI, is suited to the generation of angiographic images as well as quantitative velocimetry and rate measurements, or even to such applications as measurement of vascular compliance.

Measurement of Blood Flow Velocity PC MRI procedures are sensitive to velocity induced phase shifts in transverse magnetization. In these procedures, data are acquired first with flow-encoding gradient pulses of one polarity, followed by flow-encoding gradient pulses of the opposite polarity. The detected signals are typically processed in one of two ways. In the first approach complex differences of the acquired data are taken. Signals from stationary tissue have no velocity induced phase shift and thus cancel upon subtraction. Signals from moving spins, however, add and are preserved. In the second method, phase differences are computed. Complex difference images typically provide excellent morphological information, but only qualitative velocity information since the relationship between image pixel intensity and velocity is not linear. The phase difference images, on the other hand, provide a more quantitative presentation of velocity data because the phase shift of each pixel is directly proportional to velocity.

A significant practical shortcoming of the phase based methods is aliasing, which relates to the dynamic range of the image. The optimal SNR situation pertains when the highest measured velocity generates a phase that gives the maximum value on the gray scale image. In practice, however, the maximum velocity is not known a priori. The operator must then either choose a velocity-encoding level that encompasses a significantly larger velocity range and pay the penalty of a diminished SNR, or choose a lower value and run the risk that the maximum velocity exceeds the assigned dynamic range, in which case the measured velocity will be aliased and will appear as flow in the opposite direction. Schemes have been proposed to correct for this type of aliasing artifact, but such schemes must be used with caution to differentiate artifactual aliasing from hemodynamic effects.

Phase mapping methods have been used extensively for non-uniform (e.g., arterial) flow evaluations, which are typically performed using cardiac triggering with the acquisition of multiple phases in the cardiac cycle. Because the duration of the pulsatile cycle varies, many researchers choose to cover two cardiac cycles to ensure that velocities over the full pulsatile cycle are acquired, particularly when the aim is to measure volume flow. Alternatively, retrospective gating methods can be used. Pulsatile flow presents an added challenge since the velocity-encoding level must be chosen to encompass the peak velocity, resulting in reduced SNR in the remainder of the cardiac cycle. Schemes have been proposed to vary the velocity sensitivity through the pulsatile cycle.

Alternative spatial encoding strategies have been proposed to reduce image acquisition time. Both EPI methods and spiral scan methods have been investigated. These methods substantially reduce the time needed to acquire two-dimensional spatial information, requiring only "single-shot" excitation pulses rather than a number equal to the number of phase-encoding steps.

Fourier Velocity Encoding Phase mapping methods yield data in which the mean velocity in each image pixel is measured. Magnetization phase information can also be used to determine the velocity distribution within each pixel. This is accomplished by incorporating a cycle of phase-encoding steps (using bipolar gradients) designed to provide a phase shift proportional to spin velocity rather than spin position. An image can then be constructed in which one dimension of the image is a spatial dimension, one is a velocity dimension, and the distribution of velocities within each pixel is displayed. This method can be extended to encoding multiple spatial and/or velocity dimensions at the expense of additional scan time.

RF Pulses A key ingredient in the rapid advance made in both phase mapping and Fourier velocity-encoding methods is the use of two-dimensional selective excitation pulses. These pulses speed up data acquisition because they reduce the dimensionality of the object that must be encoded. They also reduce partial voluming and vessel overlap. Other RF pulse strategies are also attractive. Excitation by a "comb" of frequencies has been tried, which permits the simultaneous excitation of multiple and spaced parallel planes transverse to a vessel of interest.

Measurement of Wave Speed and Distensibility Increases in arterial pressure during the cardiac cycle arise from the contraction of the heart and are propagated to the vascular tree. If blood vessels were perfectly rigid, the pressure wave would be propagated at approximately the speed of sound in water. However, blood vessel walls are far from rigid, and pressure wave propagation is considerably slowed by the expansion of the vessel in response to the pressure wave. The degree of a vessel wall's flexibility is frequently referred to as its distensibility D, which can be expressed as

$$D = \frac{\Delta A}{A_0 \Delta P}\,, \tag{4.7}$$

where ΔA is the change in cross-sectional area of a vessel in response to the pressure wave, A_0 is the vessel's cross-sectional area at diastolic pressure, and ΔP is the pressure change during the cardiac cycle. Unfortunately, this definition of distensibility is not particularly useful with MR because of the difficulties in accurately measuring ΔA and ΔP.

An alternative method for the determination of distensibility is to measure the wave speed within the vessel. Distensibility is related to wave speed C by the relationship

$$D = \frac{1}{\rho C^2}\,, \tag{4.8}$$

where ρ is the density of blood.

Wave speed can be measured directly by observing the time delay of the pressure front at two stations a known distance apart. The accuracy of such a wave speed measurement is limited only by the accuracy of the distance and time measurement, and by the assumptions that (1) reflections of the pressure wave are absent and (2) the wave speed is substantially greater than the blood velocity. For many vessels these assumptions can safely be made. Consequently, detection of the onset of blood motion at each station is usually adequate.

Postprocessing As in other fields of MRI, rapid postprocessing methods are essential to permit velocity-encoding methods to be clinically useful. Conventional two-dimensional phase mapping methods require the subtraction of multiple-image data sets and the interpretation of the data in those sets. Current implementations require several minutes of data manipulation, thus preventing the use of additional scans to elucidate questionable areas.

The data contained in a velocity study can be copious and can have multiple spatial and velocity dimensions. The use of color coding, familiar to physicians from Doppler ultrasound, has been demonstrated and is potentially helpful.

Conclusions Related to MR Imaging of Blood Flow Magnetic resonance imaging can be used in a number of ways to obtain non-invasive quantitative measurements of a variety of physiological flow parameters. MRI can be used to measure aspects of blood velocity vectors in a vessel, including their spatial distribution, direction, and magnitude. Constraints on imaging time and data set size prevent the acquisition of images in which data are obtained with high resolution in three spatial dimensions, three velocity dimensions, and the temporal dimension (i.e., the cardiac cycle).

Nevertheless, useful data can be obtained in three of these dimensions with reasonable scan times. Two such applications are presented here. In the first, two spatial dimensions are obtained over the cardiac cycle. In the second, a velocity dimension is obtained with a single spatial dimension over the cardiac cycle. Additional spatial information is obtained by multiplexing the data from several stations along the vessel with a comb excitation.

Although MR blood flow imaging is coming into wide use, only a few physiological applications of MR flow measurement have been demonstrated to date. Many more exist and are likely to become important in the future.

For example, MRI can be used to monitor the response of vessel wall compliance and flow dynamics to pharmacological intervention. These MRI procedures would be useful for monitoring both fast-acting agents (such as nicotine) and slow-acting agents (such as cholesterol). Other potential applications include the correlation of changes in blood vessel morphology with changes in flow patterns, the quantitative measurement of cardiac parameters, and the analysis of model systems to advance understanding of fundamental blood flow physiology.

4.4.2 Diffusion Imaging

The effect of molecular diffusion—random displacements on a molecular scale—has been studied via nuclear magnetic resonance since the 1950s. More recently, there has been increased interest in measuring diffusion coefficients in tissue because it has been recognized that pathological processes such as stroke uniquely alter the diffusion coefficients; therefore, changes in diffusion coefficients can provide a new means for diagnosing conditions such as stroke damage. The measurement of diffusion is based on the irreversible dephasing spins experienced in the presence of magnetic field gradients.

Measurement of Diffusion Coefficients in vivo The coupling of MRI with measurements of diffusion has been developed mainly during the past decade. Despite the intrinsic sensitivity of diffusion MRI to motion artifacts, it has been demonstrated that, with care in the experimental settings and proper hardware, water molecular diffusion can be measured in vivo with MRI with fairly good accuracy and reproducibility, especially in the brain. Motion artifacts can be significantly reduced using navigator echoes or, more efficiently, EPI. With EPI, multiple images can be acquired in short time intervals, allowing greater accuracy in the measurements. Also, diffusion measurements can now be achieved in organs other than the brain, such as the kidneys or the heart.

Mapping of Diffusion Tensor Diffusion measurements may provide useful information on tissue microstructure and function, at a scale significantly smaller than the voxel size. Recently, it has been shown that the entire diffusion tensor can be determined in vivo, including in the human brain. The interest in diffusion tensor measurements, instead of simple coefficient measurements, has been spurred by the discovery that in many tissues, such as brain white matter and muscle, diffusion is highly anisotropic. This anisotropy is related to the tissue microstructure (muscle fibers, white matter fibers), which is also anisotropic, but the exact mechanisms are still unknown.

Diffusion intrinsically depends on temperature. Therefore, temperature images, or more exactly images of temperature changes, can be obtained with fairly good accuracy while the temperature of the object changes. This feature is particularly useful for monitoring interventional procedures performed within the magnet ("interventional MRI").

4.4.3 Other Tissue Parameters

One of the main strengths of MRI is that the image intensity can be made to vary significantly depending on the type of acquisition sequence. This richness of image contrast parameterization leads to many applications in which quantitative assessment of various physiological or biophysical parameters is used to distinguish disease processes. How best to use the multi-parametric nature of the imaging data to maximize diagnostic information is a research area that has by no means been exhausted. Several examples of imaging parameters that provide more than simply anatomical mapping are discussed in this section.

Relaxation Times The contrast in conventional MR images depends on the nuclear magnetic resonance relaxation times of tissues, particularly the longitudinal relaxation time T_1 and the transverse relaxation time T_2. Other relaxation times, such as T_2^*, or specific relaxation mechanisms, such as magnetization transfer, can also be determined. For each of these parameters, there are image acquisition sequences that allow direct spatial mapping of this information. Several parameters can be used in concert for maximum discrimination of tissue types as a basis for image segmentation. Quantitation of relaxation times as a means of distinguishing pathologies has not received wide acceptance, perhaps because of intrinsic biological heterogeneity or, more likely, unreliable measurement techniques for these relaxation times. However, quantitative relaxation times provide information about tissue pathology in several instances: trabecular structure, recurrent cervical carcinoma versus post-radiation fibrosis, and muscle response of exercise have all been measured with quantitative relaxation times. Relaxation times in tissues have also been shown to represent multi-component decays. This adds to the wealth of quantitative capability (and also to the unreliability of many single-component measures). Understanding the biophysical source of these various components remains a challenge, and combining them appropriately to obtain maximum diagnostic information is an ongoing research opportunity.

Oxygen Of specific clinical interest is the measurement of the oxygen provided to tissues. MRI can provide oxygen-dependent signals in the major vessels due to T_2^* differences for oxygenated blood versus oxygen-deficient blood. This is one of the mechanisms that contributes to the functional brain MR effect. Quantitation of these differences in the presence of flow presents technical challenges. More importantly, direct measures of oxygen in tissue may be achievable by looking at myoglobin or at oxygen-dependent measurements from fluorine in fluorinated hydrocarbon blood substitutes.

Strain Particularly for cardiac functional imaging, measures of cardiac strain are important indicators of heart wall function. MRI can provide such measures either using direct velocity encoding, as is done in blood flow, or using tagging techniques that lay a grid on the heart image and then look at the subsequent deformation of the grid at a later time. Both of these techniques are able to yield the basic data for strain mapping. Since strain is a tensor quantity, the eventual extraction and display of strain maps constitute a computer graphics challenge, although significant progress is being made on this frontier. A full and detailed map of cardiac strain would provide much of the information needed for a complete cardiac diagnosis. From a logistical point of view, this capability could make MRI a primary diagnostic modality for cardiac evaluation.

Very recently, it has been shown that the phase mapping methods discussed in section 4.1.1 for studying fluid dynamics can be applied to measuring the response of body tissue to harmonic mechanical excitations. For example, synchronizing an oscillating motion-encoding gradient with the actuator that induces the mechanical wave results in phase shifts in the images from which displacement maps are obtained, enabling the computation of elastic moduli.

4.4.4 Functional Brain MRI

Since its inception in 1973 MRI has evolved into one of the most powerful non-invasive techniques in diagnostic clinical medicine and biomedical research. However, MRI is used primarily as a technique for producing anatomical images. A very significant recent development, suggested by Seiji Ogawa of AT&T Bell Laboratories and applied to humans simultaneously at the University of Minnesota, the Medical College of Wisconsin, and Massachusetts General Hospital, has been the use of MRI to non-invasively map human cortical function. This new technique has been dubbed "functional magnetic resonance imaging," or fMRI, and was almost immediately reproduced by other workers in the field, using both conventional gradient echo imaging and EPI.

Contrast Mechanism The basis of fMRI lies in the fact that deoxyhemoglobin, found in red blood cells, acts as nature's own intravascular paramagnetic contrast agent. Deoxyhemoglobin, that is, hemoglobin without a bound oxygen molecule, is a paramagnetic substance. When placed in a magnetic field, a blood vessel containing deoxyhemoglobin alters the magnetic field in its vicinity. The effect increases as the concentration of deoxyhemoglobin increases, and at concentrations found in venous blood vessels a detectable local distortion of the magnetic field surrounding the vessel is produced. This distortion can affect the MR behavior of the water protons within and surrounding the vessels, an effect that manifests itself as a small but detectable change in the image intensity of appropriately acquired MR images.

Increases in neural activation within the cerebral cortex (such as are brought about by a task, stimulus, or seizure) lead to an increase in blood flow without a commensurate increase in oxygen extraction. Hence, the capillary and venous deoxyhemoglobin concentrations decrease as oxygen-rich arterial flow increases with no increased oxygen extraction. The decreased concentrations are reflected in increases in the MRI relaxation times T_2^* and T_2 and consequent increases in the signal intensity of T_2^*- and T_2-weighted MR images (the so-called blood oxygen level dependent, or BOLD, effect). Functional MRI has been successfully applied to delineate the activity in the human visual cortex using exogenous contrast agents or endogenous BOLD. The BOLD approach has already been extended to the examination of other somatosensory and cognitive tasks, such as motor movement and speech.

Imaging Techniques The further development of imaging techniques and devices will be a fruitful area for research in fMRI. For the reasons mentioned above, high-speed imaging methods are needed that are capable of covering the entire brain, to follow either an exogenous tracer or the endogenous changes that occur during activation.

A variety of techniques are capable of imaging the brain at different time scales. Echo-planar imaging (EPI) typically acquires images on the 20- to 100-ms domain. There are several fast "conventional" gradient echo techniques that image in the 1- to 60-second domain, as well as a variety of novel sequences that combine elements of both techniques, including spiral imaging and sequences known by the acronyms DUFIS and BURST. The areas of research are outlined below.

EPI was first described by Mansfield but has only recently been available commercially. In EPI, the entire MR image is encoded following a single RF excitation. The resulting imaging and reconstruction process can have a point-spread function, eddy current effects, and distortions that are substantially different from those of conventional MRI. In addition, while it is readily implementable on small-bore imaging systems, modifications to the imaging

system are required for human studies. These modifications, which remain an active area of investigation, include the use of resonant gradient systems, dedicated head gradient coils, and modified high-power linear systems. The implementation and use of modified EPI sequences on conventional scanners also represent an active area of research.

A growing set of alternatives to EPI offer different sets of advantages and are also applicable to fMRI. Spiral scanning techniques, for example, offer improved flow behavior. Many aspects of spirals present challenges, including both the efficient reconstruction and novel amplifier design for efficient implementation on whole-body systems. While both spirals and EPI attempt to cover k-space rapidly using rapidly changing gradients, other techniques, including the aforementioned BURST and a method called OUFIS, perform this coverage using rapid RF-encoding techniques. These techniques may allow fMRI without modification of hardware or the reconstruction software. The implementation and imaging issues (artifacts, point-spread function, and so on) are being actively pursued.

Finally, conventional MRI techniques are being adapted for fMRI of both exogenous and endogenous tracers. Modification of conventional techniques that improve the contrast/time relationship is important. These areas of research are particularly attractive, given the large installed base of conventional imaging systems on which such sequences would be available. Understanding and correcting motion artifacts are critical considering the small changes observed in fMRI.

Hardware Requirements Since EPI requires very fast alternating currents through the gradient coil, development of gradient coils that are torque-free, produce less acoustic noise, and/or are tailored for certain anatomic regions is currently under investigation. Also important are the bioeffects of the large dB/dt produced by these coils.

In terms of the imaging hardware as well, the requirements for successful fMRI are more stringent than those for conventional MRI. The data acquisition system must be highly synchronized, so that all pulses and data acquisition begin at precisely known ($\pm$100 ns, or better) times on each shot. This temporal coherence must be maintained for the entire length of a data acquisition run (minutes); otherwise, the interimage fluctuations add to the noise level and make brain activity harder to detect. It is important to have an fMRI data quality assurance plan operating at any facility carrying out studies over a long period of time. Problems with the scanning system that may not greatly affect normal clinical images need to be detected and corrected, since they may degrade the fMRI results much more.

Field Strength Considerations As discussed above, local field gradients induced by differences in magnetic susceptibility between the intra- and extravascular spaces induce intravoxel dephasing of water protons and result in signal contrast in the image. This intravoxel phase dispersion directly contributes to a decrease in the apparent transverse relaxation rate T_2^*. If water diffusion distances are large compared to the spatial variation in these local gradients, some of the intravoxel phase dispersion is averaged and, as a consequence, signal loss due to dephasing is diminished. Although the dephasing effect is mitigated by the diffusion-related averaging, diffusion in the presence of the field gradient is a T_2 relaxation mechanism and can lead to a decrease in T_2.

While the experimental demonstration of these effects is relatively trivial in terms of the pulse sequences required, a quantitative understanding is much more difficult to achieve. It is possible, however, to perform calculations that yield semi-quantitative predictions that can then be experimentally tested to increase understanding of the phenomenon. Based on such calculations, it is possible to predict that the field dependence of the alteration in T_2^* due to susceptibility-induced gradients around blood vessels will depend on the diffusional averaging; as such, the change in T_2^* will be a function of the motion of the protons in the brain tissue as well as the magnitude of the field gradients encountered by the tissue protons. The latter is, in turn, dependent on the size of the blood vessels and their orientation relative to the main magnetic field direction.

For a gradient-recalled echo that is sensitive to T_2^*, the signal will have between a linear and a quadratic dependence on magnetic field. These predictions are supported by data from experiments in which the signal intensity change due to visual stimulation was examined in the same individual in the same region of the brain at both 1.5 T and 4 T.

Thus, with respect to signal intensity changes (ΔS) induced by alterations in the susceptibility of blood vessels secondary to neuronal activity, the effect will increase more than linearly with magnetic field strength. However, this relationship does not guarantee that the higher field is superior for functional imaging. In order to address this question, one must consider $\Delta S/N$ where N is noise defined as the fluctuations in the signal intensity of the images obtained consecutively during the functional imaging paradigm. Contributions to this noise can come from instrument instability and physiological variables, as well as from the inherent noise in a given image. It has been shown that $\Delta S/N$ is maximum if the echo time T_E is chosen such that $T_E = T_2^*$, in which case

$$\Delta S/N \propto (S_0/N)(\Delta R_2^*/R_2^*)\,, \tag{4.9}$$

where S_0/N is the SNR at $T_E = 0$, and $\Delta R_2^* \equiv 1/T_2^*$ and R_2^* are the relaxation rates under baseline conditions and the change in $1/T_2^*$ due to neuronal activation, respectively. Each of the quantities in equation 4.9 has a potential field dependence. Each of the four quantities in this equation is known to some extent. However, the exact nature of this relationship is currently not known and is a field of substantial interest and ongoing research. Clearly, if field

strengths in excess of 1.5 T were shown to be of significant benefit to functional imaging, there would be increased interest in the development of high-field imaging instrumentation.

Processing of Functional Images In order for fMRI to become a practical neurological research and clinical investigation tool, a large number of technical issues need to be addressed, including:

1. Motion detection and compensation;
2. Characterization of temporal response patterns;
3. Characterization of physiological noise and its effects on fMRI; and
4. Display of volume functional MR images, especially in real time.

Motion Detection and Compensation: Subject motion (relative to the image-defining gradient coils) during a scanning session is a vexing problem. As brighter and darker regions move in and out of a voxel, the intensity level of that voxel will fluctuate in time, regardless of any other physiological activity. Motion coherent with the task/stimulus timing will mimic functional MR activation; motion not coherent with the task/stimulus timing will add to the noise level and make the detection of true activations more difficult. This problem is the greatest where the tissue MR signal gradient is the largest—for example, at the interface of gray matter and cerebral spinal fluid, which is also where neuronal activation occurs. Detecting and compensating for motion effects is one of the most urgent needs in functional MR imaging.

Characterization of Temporal Response Patterns: The fMRI activation signal has been shown to vary across the brain. Whether this variance represents a physiological difference in the hemodynamics of different cortical regions or a difference in the rate at which different neurons are recruited for task performance is as yet unknown. One current method for detecting activation from the functional MR time series of images is to form the correlation of each voxel's intensity time series with an ideal response time series; large correlations are deemed to signify "active" voxels. The variability in the brain's response poses questions as to which ideal function to use, how important the choice is, and how to adapt the detection methodology to allow for the different response patterns.

The correlation method can be thought of as a least squares fitting of the ideal $r(t)$ to the observed $x(t)$ via the two parameter model $x(t) \approx \alpha r(t) + \beta$, where α represents the activation magnitude and β is the mean signal level in the voxel from which $x(t)$ is drawn. Expressed in this manner, an obvious question is, How many useful parameters can be estimated from an fMRI time series? Should linear signal models be used or can nonlinear models more economically and robustly represent the patterns of signals that are seen?

Another issue is variability in response between different iterations of a mental task or stimulus. This will tend to reduce the goodness-of-fit to any low-dimensional model (and high-dimensional ideal response models are not likely to be robust). Can such changes be better allowed for using a data driven approach (e.g., principal components analysis) to extract the ideal response signal model? Can inter-iteration variability be reliably quantified and used to judge task performance and/or learning?

A related question is the characterization and detection of responses that are not "locked" to any a priori known time(s). Not all desirable neurologic investigations fit neatly into the "task A:task B" type of alternation that is the prevalent mode at present. Can fMRI be used to determine when certain neural processes are taking place, as well as where?

Characterization of Physiological Noise and Its Effects on fMRI: The dominant "noise" source in fMRI is not thermal noise in the MR receiver coils or eddy currents in the sample volume. Physiological fluctuations are several times larger and have a complex spatio-temporal structure. The causes of this "physiological noise" may include pulsing motions caused by the heartbeat, oxygenation level changes caused by respiration, and local blood flow variations on the scale of 10–20 s. To improve the techniques now used, it will be necessary to gain a physical and statistical understanding of this physiological noise.

In summary, the range of issues related to motion detection, temporal response, and physiological noise must separately and jointly be studied in a systematic fashion, probably through statistical methods. Functional MR image processing can then be adapted to extract the maximum information from the huge quantity of raw data that is generated by an fMRI scanning session.

Display of Volume Functional MR Images, Especially in Real Time: To be really useful, fMRI must be done in three dimensions, either with multislice scans or with true volume imaging. Displaying the results in a useful manner is not a trivial task, especially in real time, where the results are changing with every new acquisition. At present, two-dimensional color overlays of activated regions onto high-resolution anatomical images are the display method of choice in published papers. This technique is easily adaptable to real-time fMRI, and provides a shot-by-shot display of brain regions that are above the activation threshold. Static three-dimensional display of fMRI results is difficult enough, and no standard method has yet been settled on. Three-dimensional color rendering in real time is

certainly practical on high-end workstations now (and will become more common in the near future), but the utility of a three-dimensional rendering that is updated every few seconds is not yet clear.

An issue of special concern to EPI is the registration of the functional MR images to the high-resolution anatomical images. Distortions in the magnetic field B_0 cause distortions in the echo-planar images that are not present in the anatomical reference images (typically gathered with a fast spin echo technique). Another registration issue is the projection of fMRI activation results onto the cortical surface. One potential medical application of fMRI is brain functional mapping to aid in neurosurgical planning. For the MR-mapped functional foci to be useful to a surgeon, they must be identifiable in the surgical field. One possible approach to this problem is registration of vessels detected by MR angiography (MRA) to make surface veins visible once the craniotomy has been performed. Considering the time constraints of surgery, this would have to be done very quickly after the patient's brain has been exposed.

Safety Considerations Important to the continued development of fMRI techniques, both at high speeds and at high field, is an understanding of various bioeffects produced by the imaging. Both empirical and theoretical modeling of such effects is actively being pursued. These efforts include the measurement and modeling of nerve stimulation due to rapidly changing magnetic fields (dB/dt effect), local deposition of RF energy due to eddies produced in the body, and other effects of high magnetic field strength. Of particular relevance would be more complete models (e.g., realistic finite-element models) that include the known inhomogeneities of both the biology and the imaging systems.

Biophysical Modeling In addition to studies in the physics and engineering of imaging systems, there are research challenges in the physics and mathematics of understanding and quantifying the underlying effects of fMRI as well as modeling, in humans, the imaging techniques used to elicit such data. Unlike x-ray or nuclear techniques, the changes observed in fMRI are produced by indirect mechanisms. Observed proton relaxation rates can be affected by compartmentalized paramagnetic contrast agents including deoxygenated hemoglobin, as well as by in-flowing blood. While the observation of these changes is robust, disentangling the multiple effects and their importance (e.g., vascular vs. parenchymal changes, arterial vs. venous changes, blood volume, blood flow, and oxygenation changes, glycolytic and non-glycolytic metabolic changes) is a prime focus of current research. The question of quantification also remains: Can combinations of MR images be used to unambiguously calculate changes, in physiological units, of oxygen consumption or of blood flow?

4.4.5 Multinuclear MRI

MR Spectroscopy and Spectroscopic Imaging Magnetic resonance has the unique capability of providing in vivo information on the chemical composition of human tissues, including brain, heart, muscle, liver, and other tissues, non-invasively through a combination of imaging (MRI) and spectroscopy (magnetic resonance spectroscopy, or MRS). Spatial and temporal variations in tissue function can be evaluated by examination of the spectra and spatial distribution of ^{1}H, ^{13}C, ^{14,15}N, ^{19}F, ^{23}Na, ^{31}P, ^{39}K, or ^{35,37}Cl. The method is known as spectroscopic imaging (SI) or chemical shift imaging (CSI). Many of the challenges relating to the implementation of SI are concerned with overcoming the limitations in data interpretation due to the relatively low SNR of the data and require the application of signal processing and parameter estimation techniques such as are discussed in Chapter 13. While recent improvements in MR hardware have made possible more flexible data acquisition, the design of time-efficient sampling strategies and the development of algorithms for data analysis are areas of ongoing research. Of particular interest to physics and mathematics research are strategies for incorporating prior information by exploiting knowledge of the spatial distribution or the spectral parameters of particular metabolites.

Whereas conventional MR images have good SNR due to the relatively high receptivity and natural abundance of protons, and the high concentration of water protons in tissues, other nuclei that can be detected by nuclear magnetic resonance are of much lower abundance and have lower receptivity (e.g., ^{23}Na has 13% of the sensitivity of protons, while ^{31}P has 8% and ^{13}C has only 0.025%). Furthermore, nuclei such as ^{23}Na, ^{39}K, and ^{35,37}Cl have nuclear electric quadrupole moments. Because these ions are bound to macromolecules, with their long rotational correlation times, the electric quadrupole moments are modulated by the local electric field gradients to result in very short relaxation times and thus broad lines. This presents technical challenges but offers opportunities as well, because difference in relaxation behavior inside and outside the cell permits differentiation between intra- and extracellular environments.

Techniques for Spectroscopic Imaging with ^{31}P: The metabolites that can be studied with ^{31}P SI include inorganic phosphate, phosphocreatine, adenosine triphosphate, phosphomonoesters (which include phosphocholine and phosphoethanolamine), and phosphodiesters (which include glycerolphosphocholine and glycerolphosphoethanolamine). Localization is achieved through phase-encoding gradients alone or through a combination of volume selection and phase encoding. The phase encoding can be applied in one, two, or three dimensions, providing localization within slices, columns, or boxes. Most techniques currently use k-space encoding similar to that used for conventional MRI, followed by Fourier transforming to obtain the spatial information. Alternatively, volume selection may be achieved using a combination of gradient and RF pulses that excite either a slice, column, or box. Instead of obtaining two-

or three-dimensional arrays of intensities, as is the case with MRI, the SI data are Fourier transformed to produce a multi-dimensional array of spectra that correspond to a specific region of space. Although most studies have acquired data on a rectangular grid in k-space and reconstructed the spatial dependence of the data via a Fourier transform, alternative sampling strategies include Hadamard encoding and reconstruction and non-uniform k-space sampling. Other techniques assume prior knowledge of the spatial distribution of metabolites in order to reduce acquisition times.

Techniques for Spectroscopic Imaging with Water-Suppressed ^1H: The metabolites available for analysis by proton SI include lactate, N-acetyl aspartate, choline, creatine, GABA, glutamine, glutamate, and taurine. Other metabolites may be available depending on the SNR of the particular case.

The sequences that have been used for studying spatial variations in ^{1}H metabolites usually include water suppression, volume selection, and phase encoding. Water suppression is needed to attenuate the intense water resonance down to a level where it is possible to detect metabolites at concentrations 100 to 1000 times lower. Volume selection can be used to restrict the size of the region being studied or to reduce the contribution of spurious signals in the spectrum, such as the lipid resonances that correspond to subcutaneous fat. Two of the most commonly used techniques for acquiring ^{1}H SI data use volume selection in combination with one-, two-, or three-dimensional phase encoding. The relatively high sensitivity of ^{1}H SI has made it possible to obtain data from the brain at a spatial resolution of 1–2 cm^3 with a volume coil and 0.15–0.3 cm^3 with phased array or surface coils.

Other approaches to volume selection and water/fat suppression use spatially and frequency-selective pulses, outer volume saturation pulses, or selective inversion recovery. These techniques have been applied to obtain two-dimensional arrays of spectra from entire slices through the head with an acquisition time of about 30 minutes for three or four 15 mm slices, with a slice gap of 5 to 10 mm. While the multislice capability is useful in reducing acquisition times, the clinical evaluation of whole brain metabolism is still extremely demanding; it requires good shimming and fat suppression over the whole head, as well as acquisition of data from contiguous slices. The latter implies the use of interleaved acquisitions, which would double the acquisition time. This may prove to be feasible by means of high-speed SI techniques based on the principles of echo-planar imaging. These more-rapid data acquisition schemes demand more complex data reconstruction and thus offer opportunities for future research.

Processing and Analysis of SI Data: Experiments using two- or three-dimensional SI may produce hundreds of spectra in a single examination. Processing such data routinely would require substantial automation as well as the development of algorithms for quantifying spectral parameters in the presence of relatively low SNR. Approaches to solving these problems are discussed in Chapter 3. A further requirement for efficient interpretation of the data would be graphical tools for correlating spectra and images. Several approaches for relating metabolic and anatomical data have been proposed, including plotting arrays of spectra and superimposing corresponding grids on the MR images, making tables of the distribution of each metabolite, overlaying full MRI or edge-detected MRI data with metabolite images, and using segmentation (see section 13.1) to classify morphological tissue types in order to estimate the fractions of each tissue type corresponding to spectral voxels. For qualitative interpretation of the data, the image format is appealing but, for the experienced observer, spectral arrays are more valuable in judging the overall quality of the data, determining the metabolite levels in a specific region, and relating information about different metabolites. Thus, it is not clear at this stage which approach is best suited for visual interpretation of clinical studies and how the spectral information will ultimately be best integrated with the imaging data.

Injected Paramagnetic Contrast Agents and Hyperpolarized Noble Gases Although this report does not give emphasis to the active field of tracer chemistry as it applies to MRI, positron emission tomography (PET), and single photon emission computed tomography (SPECT), or to the mathematical aspects of the biodistribution of the agents in the body, a few important fundamental principles of biodistribution should be mentioned, along with some new developments that bear on future applications.

Injected paramagnetic contrast agents such as gadolinium or dysprosium chelates are used with MRI to track the spatial distribution over time of the agent, from which kinetic parameters such as flow or tissue perfusion can be derived using mathematical models of various degrees of complexity. These agents effect a change in the local relaxation rate $1/T_1$ that is proportional to the concentration of the contrast agent. The paramagnetic contrast agent is bound to a radionuclide tracer so that the concentration of the contrast agent in the organ or region of interest can be monitored. The sensitivity of MRI is too low for detection of the T_1 changes associated with neuroreceptor-targeted ligands or antigens in general.

Another class of contrast agents appropriate for MRI medical applications is the noble gases with spin 1/2, notably hyperpolarized ^{3}He and ^{129}Xe. As opposed to the situation with paramagnetic contrast agents, these hyperpolarized gases serve as both tracer and the source of the MRI signal. While the sensitivity of MRI is too low to detect the normal concentrations of injected nuclei, because only a few nuclei out of one million would be polarized, 10 to 20% of the nuclei are polarized with these gases, and the sensitivity increase of 10,000 to 100,000 enables their imaging. The hyperpolarized gases maintain their polarization for many minutes, or even hours if bottled in a magnetic field and stored at cryogenic temperatures.

These inert gases can be used for lung imaging and, because of their finite solubility in tissue, could potentially be used for tissue perfusion quantitation as has been done with radioactive xenon, ^{127}Xe and ^{133}Xe. Alternatively, one can utilize a known signal decay of the nuclear spin relaxation rate, $1/T_1$, to determine flow at equilibrium during constant infusion, probably through inhalation of a dilute mixture of the hyperpolarized gases.

4.4.6 Microscopic Imaging

There is no well-defined boundary between standard macroscopic imaging and microscopic imaging. A common working definition places that boundary at the best resolution attainable with the usual clinical MRI systems, perhaps about 300 μm for protons. The attainable resolution using surface coils can be less than 50 μm, so that MRI covers a volume range of about 10^6, with a corresponding range of total numbers of nuclei and hence signal strength. The potential applications range from the improved delineation of anatomical structure already resolvable by MRI to new classes of features in tissues. Examples of specialized applications that may become important include visualization of cortical layers, the precise boundaries of functional regions in the brain, and microangiography at the level of arterioles and venules. All of these applications may be augmented, or made possible, by the use of magnetic contrast agents to allow the identification of specific objects, such as blood vessels and tissues of particular types.

Under various circumstances, resolution may be limited by the achievable SNR, digitization (during acquisition or processing), bulk motion, or molecular diffusion, as well as by other problems and artifacts typically encountered in MRI.

Resolution It is common in practice to confound true image resolution with digital resolution (the dimensions of a pixel or voxel) and to emphasize pixel resolution in a thick slice. These practical working definitions can give rise to disagreements when more rigorous analyses of instrument performance and image interpretation are carried out.

There is no fundamental physical limit to resolution in MRI. The limits encountered in practice are complicated functions of the features of the apparatus, acquisition and processing methods, and object characteristics. In a given experiment, the limits may be set by the SNR in a single voxel, by digital resolution, by intrinsic line widths (homogeneous or heterogeneous broadening, by T_2 or susceptibility effects, for example), by diffusional effects, or by local or bulk motions.

Signal-to-Noise Ratios SNR may be increased by using higher magnetic fields, improved RF coil designs (including cooled coils), smaller (external, intrusive, or implanted) coils, more efficient acquisition methods, and special processing methods. When the contrast-to-noise ratio or the T_1/T_2 ratio is important, contrast agents may increase the effectiveness (the contrast attainable in a fixed time). Small objects can now be imaged with voxel resolution in the range from 100 μm^3 to 1000 μm^3, using receiver coils of the order of millimeters in diameter at magnetic fields of several tesla. It can be anticipated that regions near surfaces or implanted coils of similar dimensions will be imaged in vitro at similar resolution (about 10 μm isotropically) in the near future. Ultramicro coils (10- to 100-μm diameter, with integral preamplifier) may improve the volume resolution by about an order of magnitude because they increase the SNR by a similar factor. Because in the small-coil limit the noise arises predominantly from the coil rather than the sample, the recent development of superconducting receiver coils has afforded dramatic enhancements in sensitivity that, in turn, can be traded for improved resolution by decreasing voxel size.

Gradients As the SNR is increased, the gradient strengths required to encode image information become a problem. High-amplitude gradient coils to surround small objects are readily constructed and operated, but in vivo applications require special structures, efficient encoding methods to minimize the gradient amplitudes required, and low-duty cycles or effective cooling to avoid local heating. These considerations probably limit resolution to the order of a few tenths of a millimeter for the immediate future in human application.

Diffusion Molecular diffusion is becoming a useful source of contrast and microstructural information in studies at ordinary resolution. It plays the same role in microscopy, although unwanted diffusion damping may be a more important factor in microscopy because of the large gradients. As resolution of tens of micrometers is approached, diffusion displacements of nuclei become more important, and below a 10-μm pixel size molecular diffusion becomes the dominant consideration in experimental design.

Motion The effects of motion may be counteracted by faster image acquisition, gated acquisition (to the cardiac cycle, for example), and a variety of other tracking and processing methods already used in clinical practice and research. Postprocessing, such as realignment of images to compensate for rigid body motions, is becoming more practical as readily available computing power increases. Newer methods, such as (k,t)-space imaging, may provide a more general solution to the problem. Within some regions of the brain, motion amplitudes may be only a few hundred micrometers or less. Elsewhere in the body the problems will almost always be greater, and specific solutions may be required. It should be noted that the coordinate system is defined by the magnetic field gradients, not by

spatial location itself. Movement of gradient coils can be mistaken for object motion, and gradients that track object motion would eliminate primary motion artifacts, but are not now practical.

Future Applications of in vivo MRI Microscopy A rich variety of structural and functional features becomes visible below the millimeter level. Their visualization will usually be possible with apparatus and techniques developed for specific structures in specific locations, and for studies of a well-defined problem of function or diagnosis.

4.4.7 Research Opportunities Related to Applying Dynamic MRI

Blood Flow All MR velocity studies will benefit from hardware improvements that affect flow image quality. These include dedicated RF coils, high-performance magnetic field gradients, reduced eddy current effects, and improved spatial homogeneity, to name a few. Software improvements to enable spatially selective RF pulses are also vital. Postprocessing and display of velocity information are currently tedious and impractical. Improvements in this area must also permit the exploration of additional information in the data, such as the extraction of pressure values and shear/stress forces. Some specific research opportunities include the following:

- Development of techniques for rapid measurement of instantaneous velocity in three-dimensional space (six-dimensional problem);
- Modeling of complex flow and its implications for the vascular MR signal; and
- Development of methods for extracting parameters of physiologic relevance, including shear stress, distensibility, and turbulence.

Diffusion Imaging One field of research regards the implementation of diffusion tensor imaging (DTI) with EPI on conventional systems for clinical use. The implementation of EPI would require improvements in the gradient hardware (eddy current compensation and high-amplitude, high-slew-rate gradients), which would also benefit diffusion accuracy. Further, understanding the mechanisms underlying the diffusion values in tissues would be highly desirable. Most diffusion values are about one order of magnitude smaller in tissues than in pure water. Part of this difference can be explained by the tissue microstructure, in terms of obstacles, fibers, or membranes, but some of the differences may be artifacts. Matching MRI diffusion measurements by DTI with direct measurements using microelectrodes in tissues or animal preparations would be extremely useful. With those techniques the existence of restricted diffusion effects, requiring ultrashort diffusion times, could be demonstrated. Also the mechanisms of anisotropy in brain white matter could be better understood, as well as those involved in acute brain ischemia.

The clinical value of diffusion MRI in stroke needs more evaluation. What is the prognostic value of diffusion in terms of patient recovery? How would repeated diffusion measurements help monitor the effects of drug therapy? On the other hand, would anisotropic diffusion studies benefit management of patients with myelin disorders (e.g., brain development retardation, multiple sclerosis)? More generally, what would be the role of DTI in the evaluation of brain diseases? Still to be done are in vivo studies demonstrating the feasibility of the application of DTI to temperature imaging in interventional MRI, as well as evaluations of the effects of blood flow and tissue denaturation on the diffusion measurements.

Some specific research opportunities related to diffusion imaging include:

- Establishment and experimental validation of a molecular model for anisotropic diffusion in tissues;
- Development of improved approaches for spatially localized measurement of diffusion coefficients in vivo;
- Modeling of heat dissipation in tissue; and
- In vivo measurement of tissue fiber orientation.

Other Tissue Parameters

- Development of accurate measurement techniques for quantitative relaxation times and their interpretation in terms of clinical diagnosis.
- Development of sophisticated segmentation techniques based on multiple parametric acquisitions.
- Extraction and meaningful display of strain maps of cardiac function.

Functional Brain MRI

- Development of methods for monitoring data quality during a scanning session and for ensuring that functional activation is being observed.
- Development of new experimental protocols, especially for use with complex stimuli (e.g., visual presentations).
- Modeling and experimental verification of the biophysical contrast-to-noise mechanisms induced by neuronal activation and their dependence on magnetic field strength.
- Establishment of detailed biophysical models to understand the stimulus response permitting separation of vascular from parenchymal changes and arterial from venous changes, taking into account parameters such as blood volume, blood flow, and so on.
- Evaluation of the nature of various sources of noise (e.g., stochastic, physiological, instrument instability) and development of strategies for their minimization.

Multinuclear MRI

- Evaluation of polarized noble gases as tracers of pulmonary function and tissue perfusion (e.g., in muscle and in the brain and other organs).

Microscopic Imaging

- Exploration of the theoretical limit of spatial resolution and its dependence on key parameters, including diffusion and detection sensitivity.
- Development of improved means for monitoring of and correcting for the effects of motion, which currently limits resolution of in vivo MR microscopy.

4.4.8 Suggested Reading on Applications of Dynamic MRI

Blood Flow

1. Bryant, D.J., Payne, J.A., Firmin, D.N., et al., Measurement of flow with NMR imaging using gradient pulse and phase difference technique, *J. Comput. Asst. Tomog.* **8** (1984), 588–593.
2. Caro, C.G., Pedley, T.J., Schroter, R.C., and Seed, W.A., *The Mechanics of Circulation*, Oxford University Press, Oxford, 1978.
3. van Dijk, P., Direct cardiac NMR imaging of heart wall and blood flow velocity, *J. Comput. Asst. Tomog.* **8** (1984), 429–436.
4. Dumoulin, C.L., Souza, S.P., Hardy, C.J., and Ash, S.A., Quantitative measurement of blood flow using cylindrically localized Fourier velocity encoding, *Magn. Reson. Med.* **21** (1991), 242–250.
5. Dumoulin, C.L., Souza, S.P., Walker, M.F., and Wagle, W., Three-dimensional phase contrast angiography, *Magn. Reson. Med.* **9** (1989), 139–149.
6. Feinberg, D.A., Crooks, L.E., Sheldon, P., Hoenninger III, J., Watts, J., and Arakawa, M., Magnetic resonance imaging the velocity vector components of fluid flow, *Magn. Reson. Med.* **2** (1985), 555–566.
7. Mohiaddin, R.H., Firmin, D.N., Underwood, S.R., et al., Aortic flow wave velocity: The effect of age and disease, *Magn. Reson. Imaging* **7** (suppl. 1) (1989), 119.
8. Wedeen, V.J., Meuli, R.A., Edelman, R.R., et al., Projective imaging of pulsatile flow with magnetic resonance, *Science* **230** (1985), 946–948.

Diffusion Imaging

9. Basser, P.J., Mattiello, J., and Le Bihan, D., Estimation of the effective self-diffusion tensor from the NMR spin-echo, *J. Magn. Reson., B* **103** (1994), 247–254.

10. Le Bihan, D. (ed.), *Diffusion and Perfusion Magnetic Resonance Imaging: Applications to Functional MRI*, Raven Press, New York, 1995.

11. Le Bihan, D., Molecular diffusion nuclear magnetic resonance imaging, *Magn. Reson. Q.* **7** (1991), 1–30.

12. Reese, T.G., Weisskoff, R.M., Smith, R.N., Rosen, B.R., Dinsmore, R.E., and Wedeen, V.J., Imaging myocardial fiber architecture in vivo with magnetic resonance, *Magn. Reson. Med.* **34** (1995), 786.

13. Moseley, M.E., Cohen, Y., Kucharczyk, J., et al., Diffusion-weighted MR imaging of anisotropic water diffusion in cat central nervous system, *Radiology* **176** (1990), 439–446.

14. Turner, R., Le Bihan, D., Maier, J., Vavrek, R., Hedgers, L.K., and Pekar, J., Echo-planar imaging of intravoxel incoherent motions, *Radiology* **177** (1990), 407–414.

Other Tissue Parameters

15. Hwang, S.N., and Wehrli, F.W., The calculation of the susceptibility-induced magnetic field from 3D NMR images with applications to trabecular bone, *J. Magn. Reson., B* **109** (1995), 126.

16. Kroeker, R.M., and Henkelman, R.M., Analysis of biological NMR relaxation data with continuous distributions of relaxation times, *J. Magn. Reson.* **69** (1986), 218.

17. Labadie, C., Lee, J.-H., Vetek, G., and Springer, C.S., Jr., Relaxographic imaging, *J. Magn. Reson., B* **105** (1994), 99–112.

18. Li, K.C.P., Wright, G.A., Pelc, L.R., Dalman, R.L., Brittain, J., Wegmueller, H., Lin, D., and Song, C., In vivo verification of MR measurements of superior mesenteric vein blood oxygen saturation in a canine model, *Radiology* **194** (1995), 321–326.

19. Moore, C.C., O'Dell, W.G., McVeigh, R., et al., Calculation of three-dimensional left ventricular strains from bi-planar tagged MR images, *Magn. Reson. Imaging* **2** (1992), 165.

20. Muthupillai, R., Lomas, D.J., Rosman, P.J., Greenleaf, J.F., Manduca, A., and Ehman, R.L., Magnetic resonance elastography by direct visualization of propagating acoustic strain waves, *Science* **269** (1995), 1854.

Functional Brain MRI

21. Bandettini, P.A., Jesmanowicz, A., et al., Processing strategies for time-course data sets in functional MRI of the human brain, *Magn. Reson. Med.* **30** (1993), 161–173.

22. Bandettini, P.A., Wong, E.C., Hinks, R.S., Tikofsky, R.S., and Hyde, J.S., Time course EPI of human brain function during task activation, *Magn. Reson. Med.* **25** (1992), 390–397.

23. Belliveau, J.W., Kennedy, D.N., McKinstry, R.C., Buchbinder, B.R., Weisskoff, R.M., Cohen, M.S., Vevea, J.M., Brady, T.J., and Rosen, B.R., Functional mapping of the human visual cortex using magnetic resonance imaging, *Science* **254** (1991), 43–49.

24. Cohen, M.S., and Weisskoff, R.M., Ultra-fast imaging, *Magn. Reson. Imaging* **9** (1991), 1–37.

25. Hajnal, J.V., Myers, R., et al., Artifacts due to stimulus correlated motion in functional imaging of the brain, *Magn. Reson. Med.* **31** (1994), 283–291.

26. Kennan, R., Zhong, J., and Gore, J., Intravascular susceptibility contrast mechanisms in tissues, *Magn. Reson. Med.* **31** (1994), 9–21.

27. Kwong, K.K., Belliveau, J.W., Chesler, D.A., Goldberg, I.E., Weisskoff, R.M., Poncelet, B.P., Kennedy, D.N., Hoppel, B.E., Cohen, M.S., Turner, R., Cheng, H.M., Brady, T.J., and Rosen, B.R., Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation, *Proc. Natl. Acad. Sci. USA* **89** (1992), 5675–5679.

28. McCarthy, G., Blamire, A.M., Rothman, D.L., Gruetter, R., and Shulman, R.G., Echo-planar magnetic resonance imaging studies of frontal cortex activation during word generation in humans, *Proc. Natl. Acad. Sci. USA* **90** (1993), 4952–4956.

29. Menon, R.S., Ogawa, S., Tank, D.W., and Ugurbil, K., 4 tesla gradient-recalled echo characteristics of photic stimulation induced signal changes in the human primary visual cortex, *Magn. Reson. Med.* **30** (1993), 380–386.

30. Moonen, C.T., Liu, G., van Gelderen, P., and Sobering, G., A fast gradient-recalled MRI technique with increased sensitivity to dynamic susceptibility effects, *Magn. Reson. Med.* **26** (1992), 184–189.

31. Ogawa, S., Lee, T.M., Kay, A.R., and Tank, D.W., Brain magnetic resonance imaging with contrast dependent on blood oxygenation, *Proc. Natl. Acad. Sci. USA* **87** (1990), 9868–9872.

32. Ogawa, S., Tank, D.W., Menon, R., Ellermann, J.M., Kim, S-G., Merkle, H., and Ugurbil, K., Intrinsic signal changes accompanying sensory stimulation: Functional brain mapping with magnetic resonance imaging, *Proc. Natl. Acad. Sci. USA* **89** (1992), 5951–5955.

33. Stehling, M., Turner, R., and Mansfield, P., Echo-planar imaging: Magnetic resonance imaging in a fraction of a second, *Science* **254** (1991), 43–50.

34. Turner, R., Jezzard, P., Wen, H., Kwong, K.K., Le Bihan, D., Zeffiro, T., and Balaban, R.S., Functional mapping of the human visual cortex at 4 tesla and 1.5 tesla using deoxygenation contrast EPI, *Magn. Reson. Med.* **29** (1993), 277–279.

Multinuclear MRI

35. Albert, M.S., Cates, G.D., Driehuys, B., et al., Biological magnetic resonance imaging using laser-polarized ^{129}Xe, *Nature* **370** (1994), 199–201.

36. Bottomley, P.A., MR spectroscopy of the heart: The status and challenges, *Radiology* **91** (1994), 593–612.

37. Bottomley, P.A., Human in vivo NMR spectroscopy in diagnostic medicine: Clinical tool or research probe?, *Radiology* **170** (1989), 1–15.

38. Brown, T.R., Kincaid, B.M., and Ugurbil, K., NMR chemical shift imaging in three dimensions, *Proc. Natl. Acad. Sci. USA* **79** (1982), 3523–3526.

39. Duyn, J.H., and Moonen, C.T., Fast proton spectroscopic imaging of human brain using multiple spin-echoes, *Magn. Reson. Med.* **30** (1993), 409–14.

40. Maudsley, A.A., Hilal, S.K., Perman, W.H., and Simon, H.E., Spatially resolved high resolution spectroscopy by four-dimensional NMR, *J. Magn. Reson.* **51** (1983), 147–152.

41. Middleton, H., Black, R.D., Saam, B., et al., MR imaging with hyperpolarized ^{3}He gas, *Magn. Reson. Med.* **33** (1995), 271–275.

42. Posse, S., De Carli, C., and Le Bihan, D., Three-dimensional echo-planar spectroscopic imaging at short echo times in the human brain, *Radiology* **192** (1994), 733–738.

43. Schwarzschild, B., Inhaling hyperpolarized noble gas helps magnetic resonance imaging of lungs, *Physics Today* **48** (June 1995), 17–18.

Microscopic Imaging

44. Aguayo, J., Blackband, S., Schoeniger, J., Mattingly, M., and Hintermann, M., Nuclear magnetic resonance imaging of a single cell, *Nature* **322** (1986), 190–191.

45. Ahn, C.B., and Chu, W.C., Optimal imaging strategies for three-dimensional nuclear magnetic resonance microscopy, *J. Magn. Reson.* **94** (1991), 455–470.

46. Callaghan, P.T., *Principles of Nuclear Magnetic Resonance Microscopy*, Oxford University Press, New York, 1991.

47. Cho, Z.H., Ahn, C.B., Juh, S.C., and Lee, H.K., Nuclear magnetic resonance microscopy with 4 micrometer resolution: Theoretical study and experimental results, *Med. Phys.* **15** (1988), 815–824.

48. Chung, H., Wehrli, F.W., Williams, J.L., and Kugelmass, S.D., Relationship between NMR transverse relaxation, trabecular bone architecture and strength, *Proc. Natl. Acad. Sci. USA* **90** (1993), 10250–10254.

49. Cofer, G.P., Brown, J.M., and Johnson, G.A., In vivo magnetic resonance microscopy at 5 micrometers, *J. Magn. Reson.* **83** (1989), 608–616.

50. Early, T., Roemer, P., Mueller, O., Mogro-Campero, A., Turner, L., and Johnson, G., A high-temperature superconducting receiver for nuclear magnetic resonance microscopy, *Science* **259** (1993), 793–795.

51. Gewalt, S., Glover, G., Hedlund, L., Cofer, G., MacFall, J., and Johnson, G., MR microscopy of the rat lung using projection-reconstruction, *Magn. Reson. Med.* **29** (1993), 99–106.

52. Wu, Z., Chung, H., and Wehrli, F.W., A Bayesian approach to subvoxel tissue classification in NMR microscopic images of trabecular bone, *Magn. Reson. Med.* **31** (1994), 302–308.

53. Xiang, Q.S., and Henkelman, R.M., K-space description for the MR imaging of dynamic objects, *Magn. Reson. Med.* **29** (1993), 422–428.

Labels for Figure 4.1.

Year

$\log(\frac{1}{\text{voxel volume in mm}})$