A Practical Guide to In Vivo Spectroscopy

Dominick J.O. McIntyre

CRUK Cambridge Institute, Cambridge, UK

The overall work of which this forms a part contains highly detailed articles on specific methods, hardware, and applications of magnetic resonance (MR) spectroscopy (MRS). This article sets out to give context to the other articles by describing in practical terms how in vivo MRS is carried out, with emphasis on the layout of an MR unit, essential safety measures, and the various pieces of equipment required for a scan. The concept of a pulse sequence and the generation of its individual elements are introduced, with a brief guide to the key sequences used in vivo and recommendations on which sequence to use. Both clinical and preclinical (animal) studies are discussed.

Keywords: in vivo NMR, radiofrequency coils, instrumentation, pulse sequences, localization, safety

How to cite this article:
eMagRes, 2016, Vol 5: 995–1002. DOI 10.1002/9780470034590.emrstm1489

Introduction

A reader working through this book sequentially will have been introduced to in vivo magnetic resonance spectroscopy (MRS) at an abstract and theoretical level (see ***The Basics***; ***Magnetic Resonance Spectroscopy Instrumentation***). This article sets out to describe the daily experience of the spectroscopist working in a preclinical or clinical MRI unit. While the spectroscopist's interaction with the subject will be quite different for a human patient or volunteer compared with a mouse, the instruments and protocols are very similar, and the choice of sequence depends much more on the nucleus to be measured than the type of subject. Therefore, this article sets out a generic experimental procedure from preparation for the arrival of the subject to offline processing of the acquired data, discussing multiple options as they arise.

The Magnetic Resonance Unit

Entering the Unit

Whatever your subject may be, your first tasks as you arrive at the unit will be safety-related (see the American College of Radiology's White Paper on safety[1] and ***Safety Considerations in MRI*** for overviews; and ***Safety and Sensory Aspects of Main and Gradient Fields in MRI***; ***Safety Regulation in MR***; ***Bioeffects and Safety of Radiofrequency Electromagnetic Fields***; ***Health and Safety Aspects of Human MR Studies***; ***RF Device Safety and Compatibility***; ***Radiofrequency Fields: Interactions, Biological Effects, and Safety Issues*** for other aspects). Two principal hazards are magnetic items that are allowed too close to the magnet such that they become projectiles that injure or kill staff or patients; and implants within the body that interact with the main magnetic field or RF or gradient pulses, potentially causing burns, internal injuries, or even death.[2] The unit will generally be locked at all times. Access will be restricted to staff who are trained in magnet safety (which can be summarized as 'Don't take anything magnetic anywhere near the magnet') and who have no hazardous implants within their bodies – such items may pose as much risk in the bodies of instrument operators as in patients.

Aneurysm clips[3] and cardiac pacemakers are generally the most common and important contraindications: while exposure to strong magnetic fields can be potentially fatal and metal fragments in the eyes can blind,[4] scanning with pacemakers is possible under controlled conditions[5,6] (use of sequences with reduced specific absorption rates [SAR], and certain pacemaker functions disabled). Generally, anything within a patient's body that they were not born with must be assessed for safety and their proximity to the region of interest (ROI) assessed for potential degradation of images and spectra, before the patient may enter the magnet room. Useful screening forms may be found at the web site of the Institute for Magnetic Resonance Safety, Education, and Research[7]; this site also contains an excellent list of medical implants and their safety implications. Any patient arriving must be screened before scanning. Experimental animals generally present less of a challenge: microchips outside the area of study are usually not a problem for animals larger than mice (where the artifacts may extend across a large part of the body), but other experimental implants may not be MRI-compatible (for instance, long metallic pins for bones and devices for controlled drug infusion).

It is essential to avoid taking magnetic items close to the magnet. Modern self-screened magnets have very small stray fields, but this has the consequence that the field strength increases sharply closer to the magnet, and the resulting force on ferromagnetic objects can accelerate them into the magnet bore at dangerous speeds. You must remove all magnetic items (such as coins, watches, and bank cards with magnetic strips) before entering the stray field of the magnet.

Layout of the Unit

Unit design is driven by safety, the need to minimize noise, and institutional space limitations. Noise includes both acoustic noise emitted by the scanner and electrical noise picked up by the scanner's receiver. This last is critical for MRS, as broadband noise will quickly swamp low-amplitude metabolite signals, while coherent noise signals can mimic those of metabolites. A clinical unit will typically consist of an area for patient reception; an area for preparation of screened patients for scanning, including, removal of personal or magnetic items for safe storage; and an imaging area, subdivided into a control room, a magnet room, and a plant room as exemplified in Figure 1. The magnet room is a soundproof, electrically shielded, lockable room that fully encloses the magnet and most of its stray field (Figures 2 and 3).

Most scanners employ a superconducting magnet constructed from a coil of superconducting wire (typically NbTi) enclosed in a cryostat and immersed in a bath of liquid helium to maintain the superconducting temperature (see ***Cryogenic Magnets for Whole-Body Magnetic Resonance Systems***). A cryocooler[8] is used to prevent the liquid helium being boiled off by heat leaking into the cryostat: this device is often referred to as a *cold head*. The magnet room will also contain the receiver preamplifier, placed as close to the magnet as is practicable to minimize losses in the signal-to-noise ratio (SNR) in the cable from the RF detection coils, and electronics for connecting and switching the coils to the scanner's receiver. Other electronic components (see ***Magnetic Resonance Spectroscopy Instrumentation***) are sited in an adjacent 'plant' room, with connections into the magnet room via a filter panel that minimizes the chance of electrical noise passing through the RF shielding of the magnet room. This separate plant room both limits electronic noise pickup and simplifies maintenance by allowing use of magnetic tools without risk from the magnetic field. Sound levels in the plant room can be sufficiently high to require ear protection, so this room is also generally soundproofed.

Finally, the main computer that runs the MR scanner is located in a control room. Subjects in simple preclinical or animal studies can be prepared in the control room as confidentiality is not an issue for animals. However, those animal studies requiring sophisticated instrumentation and/or surgical preparation are best prepared in suitably equipped laboratories

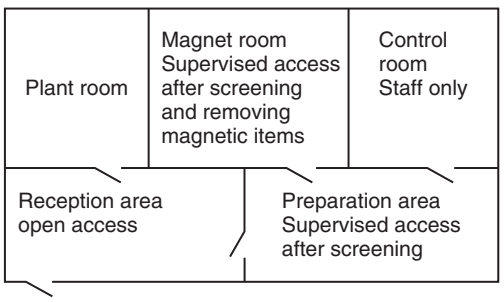

Figure 1. Example of a site layout for a clinical MRI unit, showing the stages that a patient must go through before being admitted to the magnet room where the magnetic field is potentially hazardous

Figure 2. Clinical MRI magnet. The controls on the front panel allow landmarking of a region of interest (ROI) on the patient and control of the position of the motorized bed. The shelves to the right store RF coils when not in use

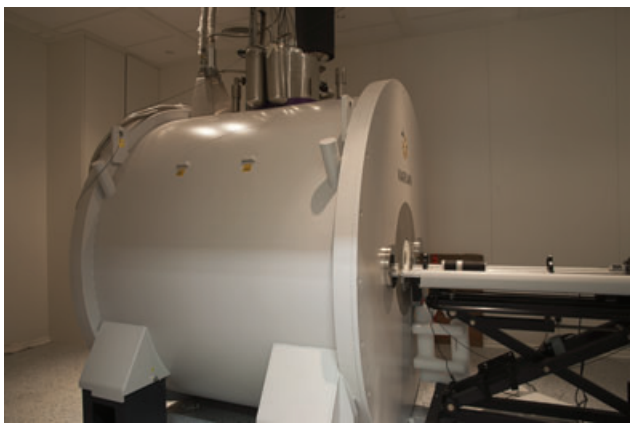

Figure 3. Pre-clinical or animal MRI/MRS system. This does not have the cosmetic casing of the clinical instrument or the motorized bed. Part of the cold head system for preventing liquid helium boil-off is visible atop the magnet

that are separated from the scanner facility and then transferred to the scanner area. Aside from such laboratory space, preclinical units generally have fewer enclosed rooms: for example, the equipment associated with smaller-bore systems may not need a dedicated plant room, and the magnet can be fitted with RF-shielded end-caps to provide an RF-shielded compartment instead of a separately shielded magnet room.

System Hardware

Typically, the operator controls the system from a computer running a standard operating system and custom software, which takes over the computer's user interface (UI) completely or which runs one or more windows in a standard UI. This will be connected to the NMR electronics via direct cabling or a dedicated Ethernet interface that allows rapid communications with no interference from other traffic or risk of interception. The operator will also have access to equipment for monitoring

the subject's physiology, and in human systems, voice contact with subjects inside the magnet bore, and typically those in the magnet room as well. Thus, only the control computer and physiological monitoring equipment are outside the plant and magnet rooms.

The plant room contains apparatus to generate waveforms for RF and gradient pulses; amplifiers for these pulses; apparatus to digitize the received signal; equipment for making the main magnetic field as uniform as possible over the ROI, a process known as shimming (see section titled 'Shimming'); the compressor for the cold head; and water cooling and air-conditioning systems for the various power supplies. Gradient and RF waveforms are generated on the control computer and transmitted to controllers within the electronics to create the MRI and MRS pulse sequences: the computer determines the exact timing of all of the pulses and the signal reception and the exact amplitudes of all waveforms for each scan. Once started (either manually by the operator or automatically based on a preset schedule), the sequence may continue using preset timings, or individual scans within the sequence may be triggered based on physiological measurements from equipment used to monitor the subject's respiration and/or heart rate, for example.

Separate digital-to-analog converter boards are used to generate control waveforms for the three gradient axes and for one or more RF pulse channels. The gradient waveforms undergo pre-emphasis processing to compensate for inductive delays and eddy current effects and then are fed to the gradient amplifiers. These can produce currents of many hundreds of amperes, which pass through gradient coils and generate the magnetic field gradients that spatially localize the MRI and MRS signals within the magnet bore. Heat generated by these currents often requires both amplifiers and coils to be water cooled. As well as the MRI/MRS gradient coils, which produce linear field gradients, high-field scanners generally include extra coils known as *high-order shims*. These generate nonlinear field gradients in the form of spherical harmonics, but unlike the MRI gradients, they are static. The shim coils can compensate for more complicated field inhomogeneities than linear shims alone, resulting in better MRS line shapes over larger volumes. Linear shims are applied using small constant currents through the gradient coils.

Frequency synthesizers or crystal oscillators are used to generate continuous RF of constant amplitude and highly stable frequency and phase. This signal is amplitude-phase- and/or frequency modulated by the RF pulse envelope to generate the RF pulse, which is then fed to the RF power amplifier to obtain the high-peak power (e.g., up to 16 kW) that is required to achieve large flip angles in a few milliseconds. The RF power amplifier output is gated off when RF pulses are not being applied, to prevent amplifier noise from contaminating the detected signal. Scanners used for proton (^{1}H) MRI and MRS only may have a single power amplifier tuned for maximum output at the ^{1}H frequency or several identical amplifiers if the scanner has multichannel parallel excitation capabilities. Scanners equipped for nonproton MRS will generally have two power amplifiers, one for ^{1}H and a broadband amplifier for other nuclei, with separate RF channels to control each. These may allow pulses at both ^{1}H and another nucleus to be played out simultaneously, enabling decoupling and nuclear Overhauser enhancement (NOE) sequences (see ***Proton Decoupling During In Vivo Whole Body Phosphorus MRS***; ***Proton Decoupling in Whole Body Carbon-13 MRS***).

The gradient cables are connected directly (or via low-pass filtering to eliminate RF noise) to the gradient coils. A single RF coil may serve for both transmission and reception, and the RF interconnections are more complicated to take account of that. For transmit/receive (T/R) coils, the same cabling must carry high-power RF pulses (on the order of 1–16 kW peak) to the coil, and the much weaker (<1 mW) excited NMR signal back to the system preamplifier. Internal switching (the T/R switch) protects the preamplifier from the high-power pulse and prevents noise from the amplifier contaminating the signal. The signal must be detected as soon as possible, and detection in <0.5 ms is often possible in practice, in simple pulse acquire and chemical shift imaging (CSI) sequences. The preamplifier is positioned as close to the magnet as possible to minimize signal attenuation in the cable. The preamplifier output returns to the receiver system within the plant room, usually via further amplification stages, where it is converted into a lower intermediate frequency by mixing with a carrier signal derived from the frequency synthesizer, and then digitized (see ***Receiver Design for MR***; ***Whole Body Magnetic Resonance Spectrometers: All-Digital Transmit/Receive Systems*** for detailed descriptions of the receiver).

RF Coils

RF coils may be transmit-only, receive-only, or fulfill both purposes. They may be volume coils, enclosing the whole body, the head (Figure 4), or a limb, or surface coils placed over a tissue ROI. Smaller coils provide higher SNR per unit volume at the expense of depth of penetration, field-of-view, and homogeneity of the transmit field and detector sensitivity.

Typically, when separate coils are used for transmission and reception, both coils will employ PIN (*p*-type intrinsic *n*-type) diode switching to shift their frequency off-resonance when not in use. This reduces cross-talk between the two coils

Figure 4. Head coil for ^{1}H MRI and MRS at 1.5 T

which would otherwise cause 'hot spots' and signal dropouts in image intensity.

Linear coils will be connected to the transmitter and/or preamplifier with a single cable. Quadrature coils[9,10] (whose circular polarization provides a $\sqrt{2}$ improvement in SNR with half the transmitter power requirement, see ***Radiofrequency Systems and Coils for MRI and MRS***; ***The Basics***) have two ports connected via a quadrature hybrid, which splits the transmitted power into in-phase and 90° phase-shifted excitation signals for the two channels during excitation and recombines the two incoming phase-shifted signals into a single channel for the preamplifier during reception. Parallel transmission systems employ two or more channels powered separately and can accommodate independent variations in phase, amplitude, and even the pulse shape for each channel, to compensate for nonuniformities in the excitation field (B_1). Parallel or phased-array receiver coils (see ***Spatial Encoding Using Multiple rf Coils: SMASH Imaging and Parallel MRI***; ***Sodickson, Daniel K.: The Many Guises of Tomography – A Personal History of Parallel Imaging***; ***Coil Array Design for Parallel Imaging: Theory and Applications***), similarly require a separate preamplifier for each channel, and the received signals are separately digitized and combined in software. On clinical systems, this is typically handled via a single multiway connector, but preclinical systems often require individual BNC (Bayonet Neill–Concelman) cable connections for each channel. Finally, ^{1}H and nonproton coils are typically connected to separate channels to allow decoupling, switching coils, and switching the MRS nucleus from the scanner without changing connections. Most coils are designed to operate at a single frequency, but dual-tuned coils may be constructed which simplify ^{1}H MRI and shimming before a nonproton MRS study (see ***Double-Tuned Surface Coils***; ***Double-Tuned Birdcage Coils: Construction and Tuning***).

It is common to require at least two probes for an MRS examination, particularly for one involving nonproton MRS. Clinical scanners generally have a whole-body transmitter coil built into the magnet inner bore and use receive-only coils appropriate for the body part under study.

The Scanning Procedure

Preparing the Subject for Scanning

Importantly, no human MRI/MRS research study can be performed in Europe or North America without an institutionally approved human study protocol, which will detail the purpose of the study, its potential hazards, subject inclusion and exclusion criteria, how informed consent will be obtained and by whom, who will perform the study, what will be done with the data, and how incidental findings will be handled. The system operator and key investigators must be familiar with these, *before* the first volunteer is scanned. The use of animals for research MRI/MRS studies will generally also require approval from an institutional animal care and use committee, which details the protocols that will be performed before, during, and after the scanning, the number and types of animals studied, and their fate. United Kingdom and Australia legislation requires licensing for the institution, each individual project, and for the personnel involved, aimed at optimizing animal welfare and minimizing the numbers used.

In either preclinical or clinical settings, the subject is admitted to the MR scanning area, prepared for the scan, and placed within the magnet with the MRI/MRS coils, and monitoring and gating that are appropriate for the study. The spectroscopist will create a new scan on the computer, enter the subject's details, and set up the sequences to be run and the coils to be used. The sequences can be loaded individually, but it is preferable to load a protocol containing all of the sequences to be used, with the correct parameters (such as number of slices, slice thickness, and echo time) already set.

In a clinical setting, the incoming subject must be consented and a checklist gone through for any exclusions. He or she will be asked to change into a gown and to remove any magnetic items and place them in a locker. The patient is then brought into the magnet room, given earplugs to protect them from acoustic noise generated by the gradient pulses, placed on the patient bed in front of the magnet, and positioned in the MRI/MRS coils. A lighting and position control system located on the magnet front panel is used to set a 'landmark' at the location of the tissue ROI in the subject, which is adjusted by moving a motorized cradle within the patient bed. The subject should be given an alarm button that can be used to attract the operator's attention during the scan: the bore contains an intercom system so that the operator and subject can converse between or during scans. If respiratory and/or cardiac gating are being used, the relevant sensors are attached and the signals checked for proper detection. Finally, the motorized cradle is advanced to insert the subject into the magnet, moving the land-marked ROI to the magnet isocenter where the field homogeneity is highest.

The above-mentioned setting assumes that clinical subjects will be awake. This is generally true, although nervous or claustrophobic patients will occasionally need mild sedation, and pediatric subjects may be sedated. Conversely, the vast majority of animal subjects will be anaesthetized before the scan. Lines for administration of any agents required during the scan will be inserted before magnet entry. Monitors for temperature, respiration, and heart rate may be connected: even if measurement of respiration is not required for gating, respiration and temperature are important for maintaining anesthesia (see ***Physiological maintenance in MRS experiments on small animals***; ***Physiological Maintenance in MRI/MRS of Large Animals*** for preclinical studies and ***Physiological Monitoring in Human MRS*** for clinical studies).

Animals will typically be placed within a handling tray, the coils positioned, the desired scan area landmarked, and the whole inserted into the magnet. If a volume coil is to be used, this will usually be locked in position in the magnet before the animal is inserted. Great care must be taken with positioning to ensure that the ROI is at the center of the magnet and the gradient coils for optimal field homogeneity and registering the target area to the coordinate frame of the localizing gradients. Unlike clinical systems, preclinical/animal systems do not normally have motorized positioning systems and repositioning can take some time.

Some coils now need to be adjusted by tuning to the exact NMR frequency and matching the impedance to the input impedance of the preamplifier, which is usually set at the characteristic impedance of the connecting cable (e.g., 50 Ω), yielding the lowest noise figure.[10] In a properly matched system, all of the power from the transmitter enters the coil, and none is reflected back to the transmitter. This minimizes the RF power required. Similarly, all of the signal is ideally received by the preamplifier, and none is reflected, which maximizes the SNR of the final spectrum. This process involves adjusting capacitors, guided by a reflection display built into the system, a portable nonmagnetic sweep generator that will show a plot of reflection in decibels as a function of frequency, an RF impedance meter, or a network analyzer. Tuning and matching with a sweep generator typically involves minimizing the reflected power on a 'scope display' and centering it at the MRS frequency. Clinical systems frequently use coils that do not require manual tuning and matching because their resonance frequency does not change significantly with patient load, and the preamplifier noise figure does not vary much over the range of loads usually encountered (e.g., 25–75 Ω).

Scanning Setup

The subject is now fully prepared for scanning, but a number of prescans are carried out before the actual scanning process begins. These include calibrating the RF power required, shimming the magnetic field to optimize homogeneity, and setting the transmitter and receiver offset frequencies. In some cases, particularly on clinical systems, these processes may be so automated that the operator will not be aware of what is happening. Alternatively, they may need to be directed to run, or it may be necessary to override automatic settings where the algorithm has not reached an acceptable state or if a different set of operating parameters are needed for a study. In addition, the operator may choose between a range of automated methods. It is therefore important for the operator to understand what is being adjusted during the prescan phase in the context of the desired experimental or clinical study objectives, even if the only possible intervention is to rerun the same process in the hope that it will work better a second time.

RF Power Calibration and B_1 *Shimming.* Some of these processes will be repeated for different scans. Conversely, for nonproton spectroscopy using surface T/R coils, power calibration may not be done routinely at all. Adiabatic RF pulses whose flip angle is constant over a wide range of RF power are often used to compensate for the highly inhomogeneous RF field strength produced by one of these coils (see ***Adiabatic RF Pulses***). Thus, a single calibration using a highly concentrated phantom solution with the coil loaded similarly to how a subject would load it, may suffice for all future experiments employing the same study protocol. Drift in the power calibration or degradation in the coil performance should in this case be assessed by routine quality assessment (QA) scans (see the following discussion).

A problem with calibration – particularly for nonproton nuclei – is that little or no signal may be visible in a single scan. In an extreme example, a time-course study of the metabolism of fluorinated drugs will yield no signal before the fluorinated agent is injected, as there is no endogenous MRS-visible fluorine in the body. Similarly, for hyperpolarized ^{13}C studies, the endogenous signal from the ROI will be far smaller than that available after substrate injection. Any calibration performed after injection must be very rapid to avoid loss of information about the initial kinetics of the injected substrate and avoid loss of hyperpolarization in those studies (see ***Pulse sequences for hyperpolarized MRS***). In such cases, it is possible to use a small, highly concentrated phantom fitted to the coil for calibration, provided that the phantom MRS signal does not overlap the resonances that are being studied experimentally. Otherwise, the experimenter's only option is to rely on prior phantom calibration, which may be a confounding factor if there are large variations in coil loading among different subjects and/or the calibration phantom.

In the absence of such issues, calibration can be carried out quickly by playing out a series of pulses with increasing amplitudes. The excitation pulses used may be nonselective, slice selective, or employ a three-dimensional (3-D) localization sequence such as point resolved spectroscopy (PRESS) to optimize the flip angle over a small ROI even when the RF homogeneity of the coil is nonuniform. For conventional RF pulses applied with a volume transmit coil, the acquired data will show a simple sinusoidal relationship between RF amplitude and signal, allowing an easy, even automatic, calibration. Adiabatic excitation pulses generate signal that increases to a plateau. Inversion pulses can be calibrated by playing out the inversion pulse at increasing amplitudes until a signal null is obtained.

When the experimenter is using a single linear coil, or a quadrature coil driven by a single channel via a quadrature hybrid, this is all that is required for RF power calibration. However, some systems now support the use of multiple independent transmit channels, which require independent calibration.[11] While it is possible to construct volume coils such as birdcage[12] (see ***Birdcage Resonators: Highly Homogeneous Radiofrequency Coils for Magnetic Resonance***; ***Birdcage Volume Coil Design***) or millipede (see ***Millipede Coils***) designs that have highly uniform RF fields when empty, when a quasiconducting sample is placed in the transmit coil the RF may no longer be uniform. This nonuniformity typically manifests as 'hot spots' or signal dropouts in the images, which can also result in excessive local power deposition (SAR) in some regions of the sample. These problems increase with the sample size and the operating frequency, as the RF wavelength in tissue decreases below the sample dimensions and anatomical structures within the sample. Thus, they are potentially much more severe for a clinical whole-body ^1H coil at 7 T where the RF wavelength in tissue is approximately 12 cm than for a head coil at 3 T or the body at 1.5 T. Driving the two ports of a birdcage with the same pulse profile, but different phases, amplitudes, and start time can improve the uniformity of excitation, reduce SAR hot spots, and the overall RF power required. The process of optimizing these parameters is termed B_1 *shimming*.[11] Transverse electromagnetic (TEM) coils (see ***TEM Body Coils***) work well at higher frequencies than birdcage coils and can be constructed with independent channels

(up to 16 at present), to provide improved homogeneity over other options. It is also possible to apply different RF pulse shapes to each channel, a process known as *transmit sensitivity encoding (SENSE)* or *parallel transmit*, thereby providing rapid excitation and localization to an ROI.

Shimming. When the subject is inserted into the magnet, susceptibility differences between the tissue and the surrounding air and within the tissue will distort the main magnetic field B_0, as dictated by Maxwell's equations. This effect is particularly important for MRS, as, if left uncompensated, inhomogeneity of <1 ppm in B_0 can distort and broaden line shapes and even cause neighboring spectral peaks to overlap. Shimming is the process of optimizing the currents in the high-order shim coils (see section titled 'Scout Imaging') to compensate for these susceptibility distortions as much as is possible (see ***Magnetic Resonance Spectroscopy Instrumentation***). Susceptibility effects scale with B_0, so higher-field systems tend to be equipped with more shim coils for providing ever higher-order corrections.

Not all tissues are shimmable to the same homogeneity: the finest line widths are available in parts of the brain remote from bone structures or air spaces, while the frontal and temporal lobes are harder to shim. In tissues like liver, that contain large amounts of paramagnetic deoxygenated blood, microscale susceptibility variations mean that very narrow lines are not attainable. Automated shimming methods include FASTMAP,[13] which analyses the phases of a series of projections through the voxel to provide an accurate shim in a few minutes; and phase-map based methods that employ gradient-echo imaging sequences to map the field offset over a slice or volume, then compute a least-squares fit of the known field profiles of the shims to the field map. It is also possible to shim, manually or automatically, on the integral of the signal recorded from either unlocalized or localized volumes. Shimming is generally carried out on the large ^{1}H water signal, even for nonproton studies.

Frequency Calibration. The shimming process will usually shift the mean field strength over the ROI slightly, and the operating frequency of the system must therefore be reset to match this. As volume selection for spectroscopy uses field gradients and frequency-selective pulses, calibration errors displace the voxel from the desired ROI (see ***Single-Voxel MR Spectroscopy***). For ^{1}H MRS, water suppression using frequency-selective pulses will be degraded by incorrect calibration. As with RF power calibration, frequency calibration is challenging when little or no endogenous signal is available. In this case, a frequency reference phantom may be used, or the required offset estimated from previous phantom studies and the experimentally observed ^{1}H water offset.

Scout Imaging

Often, spectroscopy will be performed as part of a clinical or research study in which MRS data are acquired along with MRI using multiple contrast mechanisms. While the various MRI methods fall beyond the scope of this book, we note that the MRS will at minimum require an initial set of three-plane MRI scans for localizing, followed by high-resolution MRI covering the MRS ROI, and ideally, MRI in at least two perpendicular orientations to allow accurate 3-D localization of the ROI.

MRS Setup and Acquisition

The spectroscopist must now choose the MRS localization method to be used. The first question is whether to use a single voxel or a multivoxel method. The simplest method of localization is to use nonselective pulses with a surface coil placed over the tissue of interest, which limits the signal acquired to the field of view of the coil. Generally, though, more sophisticated sequences are needed to eliminate signal outside the ROI but within the coil's sensitive volume. The principal localization sequences are PRESS,[14] stimulated echo acquisition mode (STEAM),[15] localization by adiabatic selective refocusing (LASER),[16] and image-selected in vivo spectroscopy (ISIS)[17] (see ***Single-Voxel MR Spectroscopy*** for more detail). The first three of these use double (PRESS and STEAM) or triple (LASER) spin-echoes, with slice-selective gradients applied along three orthogonal directions for each of the excitation, and refocusing pulses to acquire the signal from a volume defined by the intersection of the slices. These three sequences can in principle provide a spectrum from a single scan, although in practice multiple averages are required to improve SNR. They differ in the refocusing pulses: PRESS uses 180° pulses to provide spin-echoes. STEAM employs 90° pulses that reduce power deposition and allow for sharper slice definition and shorter echo times than PRESS, albeit at the expense of halving the SNR. With modern gradient performance allowing shorter PRESS echo times, STEAM is now used less often. LASER employs a train of adiabatic refocusing pulses: these give excellent volume definition even in the presence of the inhomogeneous B_1 provided by surface coil excitation and reduce the effects of J-coupling. However, the power deposition is higher than for PRESS and STEAM.

ISIS uses three slice-selective inversion pulses with the slice-selective gradients applied along three orthogonal directions, followed by a nonselective excitation pulse. This is repeated in an eight-step cycle with different combinations of the inversion pulses being omitted, and the signals are then combined by addition and subtraction so that all the signal from outside the ROI that lies at the intersection of the three inverted slices, cancels. This method is useful for species with short spin–spin (T_2) relaxation times, as the signal is entirely along the B_0 or $\pm Z$-axis during the preparation phase and undergoes only spin–lattice relaxation (T_1) decay. However, motion and partial saturation effects during the eight-scan cycle can result in incomplete cancellation and substantial contributions from outside the ROI to the final spectrum. Phosphorus (^{31}P) MRS is the principal application for ISIS: this is used mainly for pH measurements (see ***Measuring pHi and pHe by MRS***), studying tissue energetics (see ***^{31}P MRS of brain energetics (dimentias, stroke)***; ***MRS in the failing heart: from mice to men***; ***Muscle Studies by ^{31}P MRS***), and tumor studies.

The choice of localization sequence is largely determined by the relative sensitivity, the T_2 and/or spin–spin (J-)coupling of the resonances of interest, and how many ROIs need to be studied: for some nuclei (e.g., ^{31}P), key metabolites (such

as ATP) have relatively short T_2 values and dephasing due to J-coupling, while ^{1}H metabolites have a range of T_2s of a few milliseconds up to hundreds of milliseconds, depending on the molecular properties. Metabolites with short T_2 values require sequences with a very short delay between excitation and acquisition of signal, such as ISIS. Metabolites with longer T_2 values can tolerate longer delays, and indeed, it is possible to use spin-echo sequences with the echo time optimized for a particular J-coupled resonance of interest. The ^{1}H signal from tissue typically includes large, broad resonances of very short T_2 which are often of little clinical interest: using longer echo times can be advantageous in reducing these signals, resulting in the remaining resonances sitting on a flatter baseline, which makes quantification easier.

^{1}H-visible tissue components include water and lipid at concentrations much higher than most metabolites of interest. Water is visible at some level in all tissues at all echo time values, and its signal must therefore be suppressed to allow accurate quantification of other metabolites. This is generally done by applying one or more narrow band RF pulses centered on the water frequency followed by crusher gradients to ensure that the longitudinal water magnetization is minimal immediately before the localization sequence starts. However, it is also common practice to acquire an unsuppressed water scan that can be used to compensate eddy current effects on the spectral line shapes and/or for use as a concentration reference for metabolites acquired from the same ROIs.

Lipids have relatively short T_2 values and are often not detectable at longer echo times of 144 ms and above. In normal brain, little lipid signal is visible even at short echo times as it is mainly membrane-bound and has very short T_2. However, in pathology such as tumors, there is often an intense lipid signal associated with necrosis that can mask metabolites including lactate, which is an important marker for hypoxia.

Phase-encoding gradients may be inserted into all of the above-mentioned sequences to obtain data in parallel from multiple voxels; this is known as *spectroscopic imaging* or CSI (see ***CSI and SENSE CSI***). One-, two-, or three-dimensional schemes may be used. For a given sequence repetition time and overall acquisition time, the SNR per unit volume is the same for CSI as it is for single voxel MRS. Typically, acquisitions of this type will last upward of a few minutes, depending on the desired SNR and, for CSI, the dimensions of the CSI data set. For hyperpolarized studies, such long acquisition times are unacceptable, and these will use accelerated or single-shot spectroscopic imaging sequences if multivoxel data are required (see ***Pulse sequences for hyperpolarized MRS***); such sequences are also applicable to ^{1}H MRS studies (see ***High-Speed Spatial–Spectral Encoding with PEPSI and Spiral MRSI***).

Post-scan Procedures

Once all scans have been acquired, the subject may be removed from the magnet. If the subject is a patient, you should not discuss the scans with him or her; this is the responsibility of the clinician handling the case. If the subject is a research volunteer, and the scan reveals previously unknown pathology, you should follow your unit's procedures for incidental findings. The system operator and key investigators must be familiar with these (as well as the MRI/MRS protocol), *before* the first volunteer is scanned.

The acquired data should be backed up daily: in a clinical environment, it will probably be transferred to the local picture archiving and communications system (PACS). While these systems are excellent for image analysis, spectroscopy analysis is usually carried out in a specialized package run on an offline workstation, and the acquired data should be transferred without delay to that workstation. A number of specialized MRS fitting packages are available, which typically model spectra either by parameterized prior knowledge or using spectra of individual metabolite solutions acquired under the same experimental conditions as a basis set for fitting. LCModel is one package that uses a metabolite solution basis set (see ***Quantifying spectra in the frequency domain***; ***Advanced Spectral Quantification: Parameter Handling, Nonparametric Pattern Modeling, and Multidimensional Fitting***). Java-based magnetic resonance user interface (jMRUI) and its predecessor MRUI, and TARQUIN fit MRS data in the time domain, without previously applying a Fourier transform (see ***Time-Domain Methods for Quantifying MR Spectra***; ***Advanced Spectral Quantification: Parameter Handling, Nonparametric Pattern Modeling, and Multidimensional Fitting***).

Ideally, a daily QA scan should be performed on scanners used regularly for MRS. This involves performing a setup procedure and running a scan on a standard phantom to ensure that the RF power, SNR, and shim remain consistent from day to day: variations therefrom may allow early detection of hardware failures and pre-emptive repairs before a subject enters the scanner.

Finally, this article closes as it opened, with safety: lock up the unit before you leave!

Biographical Sketch

Dominick McIntyre obtained a PhD in NMR Microscopy from the Biochemistry Department, University of Cambridge, in 1991. After a postdoc in the MR Centre at the University of Nottingham, he has worked in in vivo MRS, DCE-MRI, and DWI applications in cancer at St George's, University of London, and more recently as MRI Instrument Manager at the Cancer Research UK Cambridge Institute.

References

1. E. Kanal, A. J. Barkovich, C. Bell, J. P. Borgstede, W. G. Bradley, J. W. Froelich, J. R. Gimbel, J. W. Gosbee, E. Kuhni-Kaminski, P. A. Larson, J. W. Lester, J. Nyenhuis, D. J. Schaefer, E. A. Sebek, J. Weinreb, B. L. Wilkoff, T. O. Woods, L. Lucey, and D. Hernandez, *J. Magn. Reson. Imaging*, 2013, **37**, 501.
2. J. F. Schenck, *J. Magn. Reson. Imaging*, 2000, **12**, 2.
3. R. P. Klucznik, D. A. Carrier, R. Pyka, and R. W. Haid, *Radiology*, 1993, **187**, 855.
4. W. M. Kelly, P. G. Paglen, J. A. Pearson, A. G. San Diego, and M. A. Soloman, *Am. J. Neuroradiol.*, 1986, **7**, 243.
5. S. Nazarian, A. Roguin, M. M. Zviman, A. C. Lardo, T. L. Dickfeld, H. Calkins, R. G. Weiss, R. D. Berger, D. A. Bluemke, and H. R. Halperin, *Circulation*, 2006, **114**, 1277.

6. S. Nazarian, R. Hansford, A. Roguin, D. Goldsher, M. M. Zviman, A. C. Lardo, B. S. Caffo, K. D. Frick, M. A. Kraut, I. R. Kamel, H. Calkins, R. D. Berger, D. A. Bluemke, and H. R. Halperin, *Ann. Intern. Med.*, 2011, **155**, 415.
7. mrisafety.com. The official site of the Institute for Magnetic Resonance Safety, Education, and Research, 2015. http://www.mrisafety.com (accessed May 26, 2015).
8. R. Radebaugh, *J. Phys. Condens. Matter*, 2009, **21**, 164219.
9. D. I. Hoult, C.-N. Chen, and V. J. Sank, *Magn. Reson. Med.*, 1984, **1**, 339.
10. D. W. McRobbie, E. A. Moore, M. J. Graves, and M. R. Prince, in MRI from Picture To Proton, 2nd edn, Cambridge University Press: Cambridge, 2007, Chap. Radiofrequency system, 175.
11. W. M. Brink, V. Gulani, and A. G. Webb, *J. Magn. Reson. Imaging*, 2015, **42**, 855.
12. C. E. Hayes, W. A. Edelstein, J. F. Schenck, O. M. Mueller, and M. Eash, *J. Magn. Reson.*, 1985, **63**, 622.
13. R. Gruetter, *Magn. Reson. Med.*, 1993, **29**, 804.
14. P. A. Bottomley, *Ann. N. Y. Acad. Sci.*, 1987, **508**, 333.
15. J. Frahm, K.-D. Merboldt, and W. Hänicke, *J. Magn. Reson.*, 1987, **72**, 502.
16. M. Garwood and L. DelaBarre, *J. Magn. Reson.*, 2001, **153**, 155.
17. R. J. Ordidge, A. Connelly, and J. A. B. Lohman, *J. Magn. Reson.*, 1986, **66**, 283.