Accurate and Efficient Localized Spectroscopy from Anatomically Matched Regions: SLOOP and Its Enhancements

Markus von Kienlin

Roche Innovation Center Basel, Basel, Switzerland

The low signal-to-noise ratio (SNR) is the single main factor limiting the spatial resolution that can be attained in nuclear magnetic resonance spectroscopy (MRS) in animals or man. This work first reviews the fundamental relationship between metabolite concentration, spatial resolution and SNR, which together define the frontier of sensitivity for MRS. It then presents SLOOP (*S*pectral *Lo*calization with *O*ptimal *P*ointspread-function) as a recipe to identify an optimal set of experimental parameters that will provide the most accurate results, in situations where geometrical a priori information about the sample is available. SLOOP, which provides robust, quantifiable spectroscopic information, has mainly been applied to study cardiac energy metabolism in healthy volunteers and in patients, and these results are briefly summarized. The quality of localized spectroscopy may be further improved through the use of nonlinear pulsed magnetic field gradients. This technology, which has been named SLOOP^N, is outlined and discussed. The text provides advice on the practical implementation of these methods with a focus on the high experimental quality required for localized MRS: even the most sophisticated postprocessing algorithm cannot fully compensate shortcomings in data acquisition.

Keywords: magnetic resonance spectroscopy, localization methods, spectroscopic imaging, spatial resolution, sensitivity, spectral localization by imaging (SLIM), pointspread function, spatial response function, acquisition weighting, higher-order gradients

How to cite this article:
eMagRes, 2015, Vol 4: 601–610. DOI 10.1002/9780470034590.emrstm1437

Introduction

In the early 1980s, magnetic resonance imaging (MRI) began invading the radiology suites of modern hospitals. At that time, the same physical phenomenon of nuclear magnetic resonance (NMR) had already been utilized for decades in analytical chemistry, for determining the constituents of liquid samples in test tubes. 'Magnetic resonance spectroscopy (MRS)' – as this method now is commonly called in the biomedical environment – was applied for the first time in living tissue in the early 1970s, to measure the essential components of phosphorous energy metabolism.[1] In the 40 years since the first biological applications, in vivo MRS has seen tremendous technical developments and has made major contributions to the elucidation of biochemical processes and pathways, in healthy and disease states.

Despite the manifold success stories of MRI and MRS, one of their major limitations is poor sensitivity, which is about six orders of magnitude lower compared to nuclear imaging methods. Fundamental physical principles, related to the specific physicochemical properties of living tissue, govern the sensitivity of in vivo MRI and MRS and limit the achievable signal-to-noise ratio (SNR). While the detected signal is proportional to the amount of the compound to be measured, the noise of in vivo measurements depends on the electric conductivity of the tissue plus the losses in the detection system.[2] Well-functioning state-of-the-art MRI/MRS instruments operate close to the natural sensitivity limit of sample-dominant noise under most practical circumstances, and except, perhaps, for hyperpolarization (see ***Hyperpolarization Methods for MRS***) no technological break-through can be expected to remedy the situation by providing substantially higher sensitivity. Typically, to be detectable by in vivo MRS within an acceptable experimental scan-time, a tissue constituent must be present at a concentration in the millimolar range, in a volume of at least several hundred microliters.

The availability of methods to control the spatial origin of MRS signals, which like MRI are based on radiofrequency and magnetic field gradient pulses, is a key asset of in vivo MRS. They enable its use for 'noninvasive biopsy' in which the chemical constituents are measured within a well-defined region in the body (e.g., in a specific organ or lesion), remotely and without injury. The SLOOP (*S*pectral *Lo*calization with *O*ptimal *P*ointspread-function) method, which is reviewed in this article, is an approach aimed at achieving the highest-possible accuracy for spatially localized MRS. The fundamental concept is to optimally acquire the potential signal by exploiting a priori information about the anatomical shape of the region-of-interest (ROI). This article provides: (i) a recipe to determine the best experimental parameters for the MRS measurement

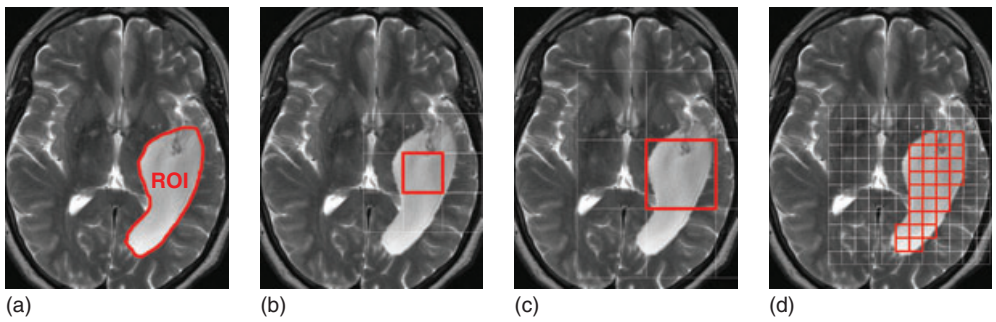

Figure 1. Example of a typical situation for localized MRS studies: (a) a lesion is detected on anatomical images and further diagnostic information is wanted. The goal of the MRS study is to measure with the highest possible accuracy the constituents in this 'region-of-interest (ROI).' The operator can position either (b) a small voxel well within the lesion to minimize contamination from adjacent tissue or (c) a large voxel centered on the lesion to maximize the signal-to-noise ratio; it is also possible (d) to accumulate the signals from many small voxels positioned within the ROI

using SLOOP; (ii) the formulae for reconstructing the localized MRS signals, taking into account the geometric shape of the ROI and other experimental parameters; and (iii) quality criteria for documenting the SNR efficiency and quality of localization for a specific study. Selected examples of applications are presented, advice on practical implementation is given, and possible technical developments are outlined.

The Spatial Frontier of In Vivo MR Spectroscopy

The fundamental goal of localized in vivo MRS is to accurately measure the chemical constituents in a target area, or 'ROI', within the living body. This task is exemplified in Figure 1(a). Assume that some lesion has been detected in the brain, but its exact nature still needs to be verified. MRS in principle can noninvasively biopsy it and determine its chemical constituents, to potentially arrive at a more precise diagnosis and/or aid in designing or monitoring an appropriate medical treatment.

There are many different possible experimental approaches for controlling the spatial origin of the MRS signal by employing sequences of radiofrequency and gradient pulses. A common property of most methods is that the detected signal is confined to a more-or-less cubic-shaped region, or 'voxel', which can be readily positioned in the ROI as depicted in Figure 1(b). The spectroscopic signal detected from such an experiment largely originates from that voxel and thus will reflect the chemical composition of that area. For a given experimental setup, the signal amplitude detected from the voxel is directly proportional to the concentration C of the substance, to the volume V of the selected voxel, and to the total duration of the measurement, T_{tot} (equivalent to the number of accumulated transients). Noise being uncorrelated, increases with the square root of the total experimental time. The SNR of the measurement thus can be written:

$$\mathrm{SNR} \sim V \cdot C \cdot \sqrt{T_{tot}} \qquad (1)$$

For an MRI or MRS measurement to be of practical use, its SNR needs to be above a certain level, to ensure accuracy and reliability. The dependency of the SNR on concentration, voxel volume, and experimental duration described by equation (1)

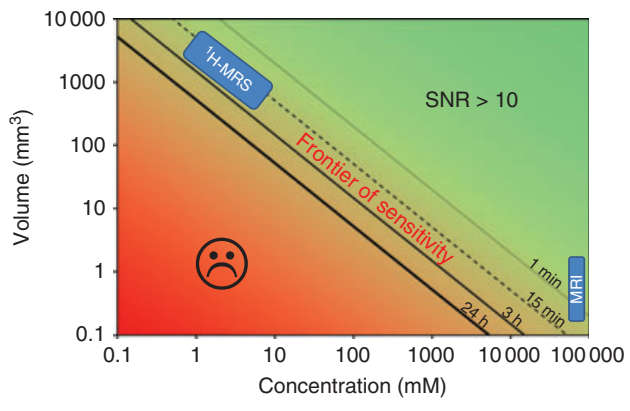

Figure 2. Schematic representation of the signal-to-noise ratio (SNR) of MR studies in living tissue, as function of concentration and voxel volume (equivalent to spatial resolution). Only in the green domain is the SNR sufficiently high to allow the generation of useful MR results. Prolonging the experimental duration to hours or days can push the 'frontier of sensitivity' to somewhat higher spatial resolutions, but concentrations in the micromolar range and below are clearly outside the reach of in vivo MR methods

is schematically depicted in Figure 2. As 'voxel volume' is equivalent to spatial resolution, Figure 2 is a visualization of the spatial limits of MR based on its inherent sensitivity. The two blue with blue labeled regimes indicate the operational range of typical methods in this landscape. Conventional MRI detects the resonance of hydrogens in tissue water present at a concentration of about 80 M, which permits useful imaging with resolution limits in the millimeter range at experimental durations of minutes (blue area in the lower right). Cerebral ^{1}H-MRS, in contrast, detects metabolites present in millimolar concentrations: to achieve a useful SNR, the resolution is limited to milliliter volumes ($\sim$1000 mm^3) for experiments requiring up to 1 h (upper left).

Note that this schematic is just an order-of-magnitude representation: the SNR depends on many other parameters including magnetic field strength, detector, sample properties, measurement protocol, etc. Nevertheless, these other parameters do not dramatically affect the fundamental resolution

limits imposed by the inherent sensitivity of the MRS experiment, which are determined by the electrical properties of the living sample on the noise side, and the spin density that generates the signal.[2] In addition, for each line on the graph, the product of the compound's concentration and voxel volume is constant. This means that the line represents a fixed absolute amount of compound in the micromole range (10^{-6} mol). MR is thus about six orders of magnitude less sensitive than nuclear imaging modalities that are capable of detecting compounds down to the picomole range (10^{-12} mol).

Returning to Figure 1, the operator of the MRI/MRS instrument has to make a choice: he can choose a rather small voxel that fits completely within the target region (Figure 1b) to provide a result that is not contaminated by signals from the surrounding tissue. However, owing to equation (1), the choice of a small voxel volume leads to a low SNR, which may compromise the accuracy of quantitation of the detected metabolites. The operator could alternatively choose a larger volume providing a better SNR: the voxel in Figure 1(c) has a threefold larger area so the detected signal will have proportionally higher SNR. The accuracy in detecting the chemical composition within the lesion is now nevertheless compromised by the inseparable signal contamination originating from tissue outside the lesion.

To find a better compromise between sensitivity and signal contamination, one may also propose using chemical shift imaging (CSI; see ***CSI and SENSE CSI***) with high spatial resolution, then sum the signals from all the voxels within the ROI (Figure 1d). While this may be a valuable approach in high *SNR* situations, there is a severe penalty in sensitivity to pay, as demonstrated in the following. In a first thought experiment, let the signal of the whole ROI be S_1, and the standard deviation (SD) of the noise be σ, for a fixed total scan time with N scans. The theoretically optimal SNR_1 of the whole ROI thus is:

$$\mathrm{SNR}_1 = \frac{S_1}{\sigma} \tag{2}$$

In a second experiment, let us now divide the ROI into N subvoxels of equal size, which are spatially encoded using a CSI sequence with N phase-encoding steps, so that the total scan time is the same as the first experiment (Figure 1d). The signal in each voxel is $\sim S_1/N$, and the summed signal S_N from all N voxels is:

$$S_N = \sum_{i=1}^{N} \frac{S_i}{N} = S_1 \tag{3}$$

just as in the first experiment. The SD of the noise of the spectrum in each subvoxel is again σ, as the noise is not affected by the phase encoding. When summing the signals from the N subvoxels, the noise adds incoherently with $\sqrt{N}$, and the SNR_2 from the whole ROI for the second experiment is:

$$\mathrm{SNR}_2 = \frac{S_1}{\sigma \cdot \sqrt{N}} = \frac{\mathrm{SNR}_1}{\sqrt{N}} \tag{4}$$

Comparing equations (2) and (4), the penalty in sensitivity to be paid by summing N voxels in the second experiment thus amounts to $\sqrt{N}$, which might only be acceptable in high *SNR* situations.

This example also substantiates the conclusion that the sensitivity of an MRS measurement is predetermined during the experiment, at the time the signals are acquired: if a measurement is run with unsuitable parameters – in particular, with too high a spatial resolution – the loss in sensitivity cannot be regained by postprocessing, and the result will always remain suboptimal. The question addressed by the SLOOP technique therefore is: what is the optimal data acquisition strategy to obtain the MR signal from a noncubical-shaped ROI, and how can signal contamination be minimized and sensitivity be simultaneously optimized?

Localization of Anatomically Matched Regions

It is a common principle of experimental physics to utilize all information known about a system with the aim of increasing the accuracy of the result of a measurement. This a priori information, together with the experimental parameters of the measurement protocol, is used to construct a simulated model of the system. This mathematical model is then utilized during data processing to determine with higher accuracy the desired parameters from the measured signals. The better the model reliably captures all of the known properties of the experimental situation, the higher the potential gain in accuracy.

This concept of using a priori information has been proposed for processing localized in vivo spectra from noncubically shaped regions in an ingenious idea from Hu *et al.*,[3] which they named 'spectral localization by imaging (SLIM)'. The basic idea is to compute a model of the spectroscopic experiment based on information obtained from an anatomical image. On this image, all ROIs, and all other regions that make contributions to the detected spectroscopic signal, are identified by a segmentation process, wherein each point in space is assigned to one out of a list of N specific compartments. The guiding principle in this segmentation procedure is to conglomerate all subvoxels or 'pixels' within a homogeneous region, i.e., all pixels generating a more-or-less identical spectroscopic signal, into one compartment. This compartment information is then used in computing a model of the detected signal in a phase-encoded spectroscopic (CSI) experiment. Let c_n be the (unknown) spectroscopic signal of the nth compartment in the absence of phase encoding and M the number of phase-encoding steps ($M > N$). The signal p_m detected for the mth phase-encoding step can then be written as the sum of the phase-encoded signal from all N compartments:

$$p_m = \sum_{n=1}^{N} g_{mn} \cdot c_n \tag{5}$$

In this equation, the variable g_{mn} captures the effects of phase encoding on the nth compartment:

$$g_{mn} = \int_{\text{compartment } n} e^{-\mathrm{i}\cdot\vec{k}_m\cdot\vec{r}} \cdot d^3\vec{r} \tag{6}$$

with $\vec{k}_m$ as the mth phase-encoding gradient. [The effect of inhomogeneous RF B_1 field transmission or detection coils can

also be captured[4] in equation (6), but this has been omitted in this article for the sake of simplicity.] All variables g_{mn} can be written as the $M \times N$ matrix **G**. This is the mathematical model of the SLOOP experiment, relating the signal from each compartment to the available information from the experiment: the phase-encoding gradients from the MRS acquisition and the anatomical distribution of the subject under study. As detailed in Ref. 3, the spectroscopic signals c_n from the N compartments can be reconstructed by computing the pseudoinverse of the matrix **G**, thereby obtaining the spectra of anatomically matched ROIs.

Some caution is warranted in the numerical implementation of these equations, to appropriately account for intravoxel dephasing when evaluating equation (6). Phantom studies have demonstrated that it is not sufficiently accurate to compute just the central phase change $\vec{k}_m \cdot \vec{r}$ for each pixel on the anatomical image (at location $\vec{r}$ in laboratory space). The intravoxel phase dispersal along each phase encode direction and the subsequent amplitude change need to be accounted for by integrating the effect in the interval $[r - {}^{\Delta r}/_2, r + {}^{\Delta r}/_2]$ over each of the three orthogonal voxel edges of length Δr. The precise coordinates of the pixels in the anatomical image also need to be known, in order for the model to accurately reflect experimental conditions. Typical implementation of the fast Fourier transform (FFT) algorithm, for example, shifts the image to the left by half a pixel dimension in each direction so that the zero frequency falls exactly in the center of one pixel. The exact coordinates of the left and right edges of an image about the isocenter are thus $\left[\frac{-\mathrm{FOV}-\Delta r}{2}, \frac{\mathrm{FOV}-\Delta r}{2}\right]$ rather than $\left[\frac{-\mathrm{FOV}}{2}, \frac{\mathrm{FOV}}{2}\right]$, with FOV and Δr the field-of-view and the digital resolution of the image in each dimension, respectively.

Spatial Response Function: The Optimal Experiment

The SLIM approach to solving equation (5) for the desired local signals c_n, by computing the pseudoinverse of the matrix **G**, will return a result in most practical circumstances, but in degenerate situations there may be no mathematical solution. This section addresses three fundamental questions about the quality of the result thus obtained: (i) What is the localization quality, i.e., how much contamination from surrounding tissue may have leaked through? (ii) What is the SNR efficiency of the approach? (iii) What are the optimal experimental parameters, i.e., which set of phase-encoding gradients will minimize spatial contamination and maximize SNR? The SLOOP method[5] addresses these questions; it provides quantitative criteria to evaluate the quality of localization and SNR efficiency, and a strategy to identify the best data acquisition parameters for a specific experimental situation.

The appropriate tool to analyze the properties of a localized spectroscopic experiment is to compute its *spatial response function* (SRF), based on the anatomical segmentation, the SLIM reconstruction coefficients, and the phase-encoding gradients [equation (10) in Ref. 5]. The original SLOOP publication[5] misappropriated the term *pointspread function* (PSF) for what should have been called SRF: the two have quite different physical meanings, although their mathematical functions may often look very similar. It is a frequent inaccuracy to equate 'laboratory space $\vec{r}$' with 'image space $\hat{r}$', but one should realize that – together with $\vec{k}$-space – there are indeed three relevant spaces in MRI! The laboratory space is the universal space we work in, with its three continuous dimensions x, y, and z. Image space in contrast is a space with one, two, or three discrete dimensions, stored somewhere in a computer memory or visualized on a computer screen: it is easily confounded with laboratory space as the former is a more or less accurate representation or 'image' of the latter. Both the PSF and the SRF establish the relationship between laboratory space and image space, but from reciprocal perspectives. The PSF is the mathematical function that describes how the signal from a singular point in laboratory space propagates across image space. Given a point $\vec{r}$ in laboratory space, the PSF is thus a discrete function with a value for each pixel in the image. The SRF on the other hand is the function that – for a singular pixel in the image – describes the weighted contribution to its signal arising from each point in laboratory space. Given a pixel $\hat{r}$ in image space, the SRF is thus a continuous function in laboratory space.

SLOOP is an exemplary case where the PSF and SRF are not alike. In an experiment with N compartments defined on an anatomical image, there is a PSF for any point in laboratory space, and any PSF has exactly N discrete values. On the other side, there are exactly N SRFs, one for each compartment, and each SRF is a continuous function over laboratory space.

The SRF is the appropriate tool to analyze the properties of a localized MRS experiment. It depicts the spatial origin of all contributions to the spectrum, both in amplitude and phase. One particular feature inherent to the SLIM reconstruction method is that the integral of the SRF over the ROI is exactly 1, while the integral over each other compartment defined on the segmented anatomical image is exactly zero. A first application example is shown in Figure 3: the three compartments 'cortex', 'papilla', and 'renal pelvis' have been outlined on the anatomical image of a rabbit kidney. The SRF for each compartment (obtained after a SLOOP optimization of the experimental parameters, as described below) is visualized in panel b: the amplitude of the SRF (left column) is high within each ROI but low on the outside, whereas its phase (right column) is uniformly close to zero within each ROI but dispersed over the full range elsewhere.

In an ideal experiment, the amplitude of the SRF is at some constant level within the ROI, and zero across all other compartments. The resulting localized spectrum then reflects the mean of the metabolites within the ROI with equally weighted contributions from each point in the ROI, whereas all signals from outside the ROI are efficiently suppressed. In real life, this ideal situation cannot be achieved, because phase encoding with linear magnetic field gradients cannot induce local phase discontinuities. The suppression of nonzero amplitude signals originating from outside the ROI is instead achieved compartment by compartment, through phase cancellation of the contributions within each compartment. Put simply: each positive signal is compensated by a negative signal of equal amplitude.

This phase cancellation, however, can only work if the signal from each point in the compartment has equal strength. Consequently, the SLIM reconstruction method is vulnerable

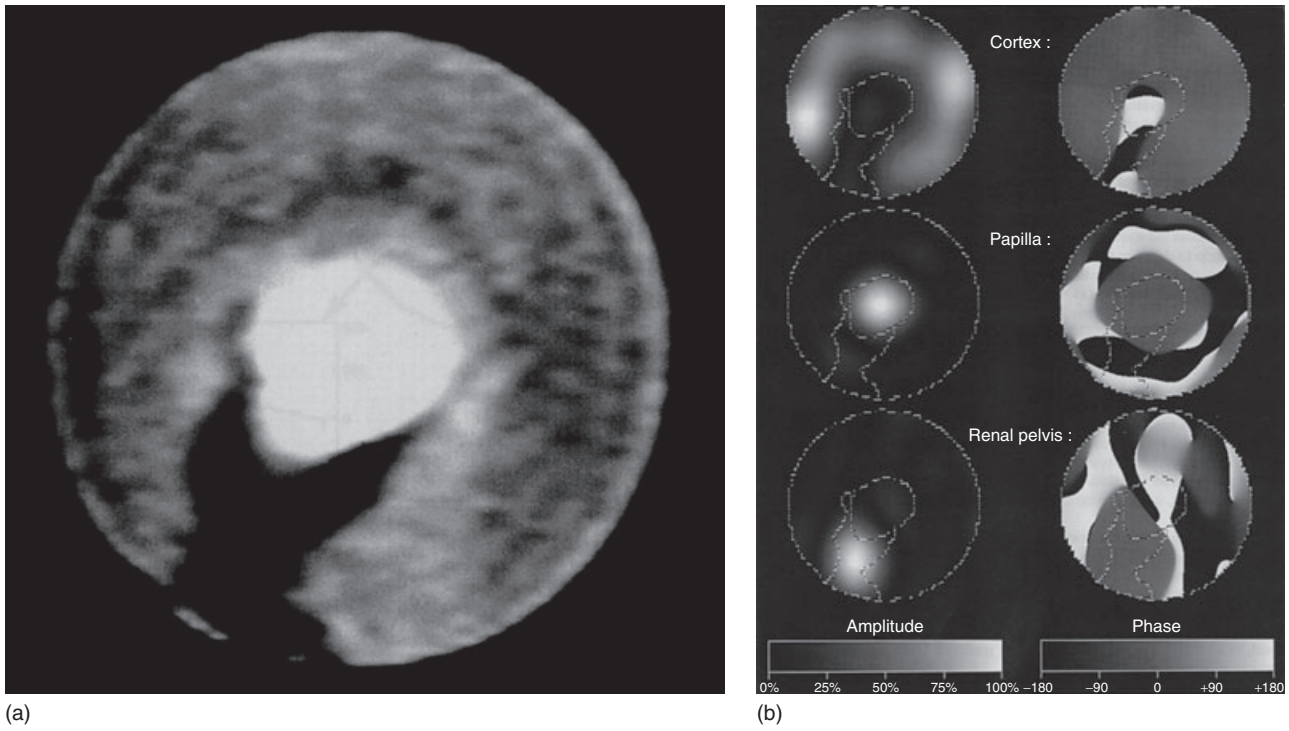

Figure 3. Example for localized spectroscopy from anatomically matched compartments: (a) on the structural image of an excised rabbit kidney, three major anatomical compartments can be readily identified, i.e., the cortex, the papilla, and the renal pelvis. Panel (b) shows the SRFs obtained for each of the compartments in a SLOOP experiment: the amplitude of the SRF matches quite well the outline of each ROI, while its phase is homogeneous within but greatly dispersed outside the ROI. (Reprinted from *Journal of Magnetic Resonance* 94: 268–287; M. von Kienlin and R. Mejia: "Spectral Localization with Optimal Pointspread Function" (1991), with permission from Elsevier)

to any inhomogeneities in the distribution of metabolites that are not taken into account in the segmentation process. Violations of the assumption of homogeneous compartments can lead to spatial contamination with signal contributions from other regions. In practice, one indication for such spatial contamination is aberrant phasing behavior of the localized spectra. If a whole set of localized SLOOP spectra cannot be easily phase corrected into pure absorption line spectra with a single set of zero- and first-order phase correction parameters, then poor localization may be the culprit.

A localization method that generates an SRF with low amplitude outside the ROI is less prone to spatial signal contamination caused by local inhomogeneities. A quantitative criterion L for comparing the localization quality of SLOOP experiments, based on the amplitude of the SRF outside the nth ROI, can thus be defined as:

$$L_n = 100\% \cdot \int_{\text{outside ROI}_n} |\text{SRF}_n(\vec{r})|\, \mathrm{d}\vec{r} \qquad (7)$$

The smaller L is, the better the SRF matches the anatomical shape of the region of interest, the less is the experiment prone to spatial contamination, and the more reliable is the result. At the same time, the SNR efficiency of the experiment should be as close to optimum as possible. This efficiency, E, can be calculated using the formula derived in Ref. 5, with V_n the volume of the nth ROI, M the number of phase-encoding steps, NA_m and NA_{tot} the number of acquisitions for the mth phase-encoding step or the total experiment, respectively, and h_{nm} the reconstruction coefficients obtained from the SLIM reconstruction process:

$$E_n = \frac{100\%}{V_n \cdot \sqrt{\text{NA}_{\text{tot}}} \cdot \sqrt{\sum_{m=1}^{M} (\text{NA}_m \cdot |h_{nm}|^2)}} \qquad (8)$$

The ideal value for E is 100%, which would correspond to a situation where all signal available in the ROI is detected and none is lost to phase dispersal. SNR efficiency, however, is in conflict with localization quality. Phase-encoding gradients are required to achieve localization, and stronger gradients produce better spatial resolution, but also lead to higher phase dispersal and consequently lower signal amplitudes. The optimal experiment has to seek the best compromise between simultaneously minimizing spatial contamination and maximizing SNR efficiency. No analytical solution is presently known for identifying the optimal experimental parameters that yield the best accuracy. However, a numerical optimization algorithm was used for this purpose in the original SLOOP implementation. In a computationally brute force approach, the strength of each phase-encoding gradient was set, and the quality criteria L and E were calculated. In an iterative process, the phase-encoding gradients were then modified until no further improvement of

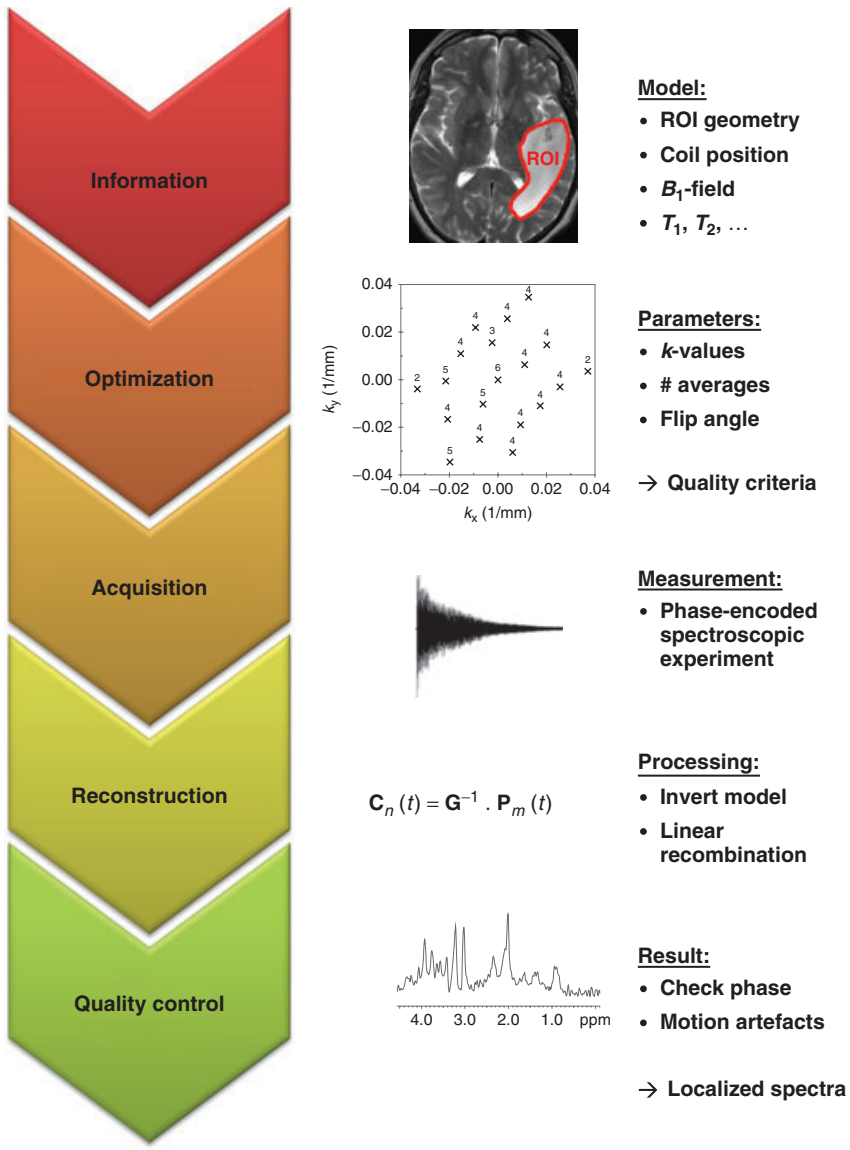

Figure 4. The five steps in the SLOOP process for accurate quantitative localized MRS. First, all available information about the experimental setup is gathered in a mathematical model, which captures the particular geometries of the ROIs. The optimal experimental parameters for that setup are then determined in a numerical optimization algorithm, which also computes measures of the quality of localization and signal-to-noise efficiency. With these parameters, the spectroscopic measurement is conducted. Local spectra are reconstructed by computing the inverse of the mathematical model and applying these coefficients to the phase-encoded raw data. Results should always be checked for signs of measurement artifacts, in particular for phase anomalies.

the quality criterion could be obtained, to yield the optimal gradient set for the experiment.

The full process of obtaining localized spectra with the SLOOP approach is depicted in Figure 4. It consists of these essential steps: all available a priori information of the studied object and the experimental setup are captured in a mathematical model. In particular, the anatomical images are segmented, assigning every signal-contributing pixel to one of a list of compartments. If coils with inhomogeneous radiofrequency fields are used for transmission or reception, they are also captured in the model.[4] A numerical algorithm then searches the best experimental parameters for this given experiment. In particular, it determines a set of *k*-values for the phase-encoding experiment, which yields the best compromise between spatial resolution and signal-to-noise efficiency. If needed, further parameters such as excitation flip angle or the number of averages per phase-encoding step can also be included in the optimization. Finally, the actual measurement is conducted with the experimental parameters computed by this optimization algorithm. To reconstruct the localized spectra, the mathematical model that describes the measured signals is inverted; the local signals then can be computed by

a linear combination of the measured signals, weighted by the reconstruction coefficients of the model inversion.

Biological Applications in Kidney, Eye, and Heart

The SLOOP method for localized spectroscopy achieves its full potential in situations where meaningful a priori information on compartments with different metabolic content can be derived from structural images. The first biological sample in which this was successfully demonstrated was an excised rabbit kidney.[5] Figure 5 shows the two main resonances that can be detected by ^{1}H-MRS: the signal of trimethylamine (at 3.1 ppm) and a strong lipid signal (around 1.3 ppm); lipids are only found in the renal pelvis. The excellent localization quality of the measurement can be judged from the absence of the strong lipid signal in the spectra originating in the renal cortex or the papilla.

Another initial study conducted in the rabbit eye perhaps had more biomedical relevance: ^{19}F-spectroscopy was used to measure the enzymatic activity of aldose reductase in the lens.[6] The eye again is readily amenable to SLOOP, as the various signal-contributing compartments can be easily segmented on anatomical images. This method can be used to measure the efficacy of aldose reductase inhibitors, which have the potential to prevent cataract formation in diabetic patients.

The major subsequent application for SLOOP has been in human cardiac phosphorous (^{31}P) spectroscopy. SLOOP was initially implemented and validated[4] for this application on a clinical MR scanner by Ralf Löffler in collaboration with Siemens Medical Systems. This experiment is illustrated in Figure 6: on each plane of the three-dimensional anatomical image, the main compartments such as 'chest wall', 'liver', 'left ventricle', or 'left ventricular blood' are identified (Figure 6a). Figure 6(b) and (c) show two planes of the 3-D SRF of the optimized experiment for the left ventricular wall: the origin of the signal is well contained within the ROI and with uniform phase; the decreasing amplitude toward the dorsal portion of the heart is due to the decreasing sensitivity of the surface coil. The localized spectra originating in the various anatomical compartments show the expected patterns confirming the high quality of localization. In particular, no phosphocreatine (PCr) signal is seen in the spectrum assigned to the liver, the PCr-to-adenosine triphosphate (ATP) ratios in the chest wall and the heart correspond to the expected values, and no PCr nor ATP is detected in the blood compartments.

This experimental setup for optimized ^{31}P-MRS of the human heart was then further refined to demonstrate accurate quantification of absolute high-energy phosphate concentrations in the left ventricular wall.[7] These successful pilot studies finally lead to a series of clinical studies in healthy volunteers and patients. Köstler *et al.*[8] found no difference in the concentration of high-energy phosphates in the left ventricular wall between male and female volunteers, but a significant reduction was found in volunteers older than 40 years. A study of left ventricular ^{31}P metabolites in patients with hypertrophied or failing myocardium by Beer *et al.*[9] demonstrated the importance of absolute quantification: both PCr and ATP were found to be significantly reduced in patients with dilated cardiomyopathy, with reductions correlating with measures of declining function, whereas changes in the relative PCr/ATP ratio underestimated the pathologic changes in absolute high-energy phosphate levels. Recently, the same group reported that the SLOOP approach can not only detect the impaired cardiac high-energy phosphate metabolism in patients with Fabry disease, but also the improvement in levels after treatment with enzyme replacement therapy.[10]

Beyond *k*-Space: Nonlinear Gradients

As already briefly mentioned in the section titled 'Spatial Response Function: The Optimal Experiment', the linear gradients used in conventional MR instruments can only induce linear phase gradients across the object, which leads to a natural limitation both for avoiding spatial contamination and for achieving the best sensitivity. A higher quality of both localization and SNR efficiency might be achievable if more pronounced phase discontinuities were inducible at the compartment boundaries. In principle, this is possible to some extent with gradient coils that generate nonlinear field dependencies in space. Rather than designing and building dedicated gradient coils for each experimental situation, a more practical approach is to utilize the shim coils installed on most MR systems to homogenize the static magnetic field. These shim coils are typically built to generate fields decomposed in spherical harmonics, which can efficiently match any specific field dependence (see ***Magnetic Resonance Spectroscopy Instrumentation***). The electronics of the scanner need to be modified to allow controlled pulsing into the shim coils, and precautions may be necessary to prevent or compensate for eddy currents (which may already exist on some state-of-the-art MR instruments). In any case, it has been shown that applying nonlinear phase encoding with pulsed currents to the channels of a shim coil system can improve the quality of the SLOOP-localized spectroscopic experiment. This approach has been named 'SLOOP^{N}' to indicate the use of higher-order gradients.[11]

The *k*-space formalism, which is very convenient for mathematically describing imaging methods as well as SLOOP [equation (6)], is unsuitable for analysis of the effects of nonlinear phase-encoding gradients – in this sense SLOOP^{N} reaches beyond *k*-space. A more general formalism can capture the local effect of all nonlinear gradient coils combined. The dynamic local field generated at time t during the mth phase-encoding step can be expressed as:

$$B_m(t,\vec{r}) = \sum_s I_{m,s}(t) \cdot B_s(\vec{r}) \tag{9}$$

with I_s the current in the sth coil, which generates the local field $B_s(\vec{r})$ per unit current. If phase-encoding currents are pulsed during the time interval $[t_1, t_2]$, the local phase change $\varphi_m(\vec{r})$ in the mth phase-encoding step is given by:

$$\varphi_m(\vec{r}) = \gamma \cdot \int_{t_1}^{t_2} B_m(t',\vec{r})\, \mathrm{d}t' \tag{10}$$

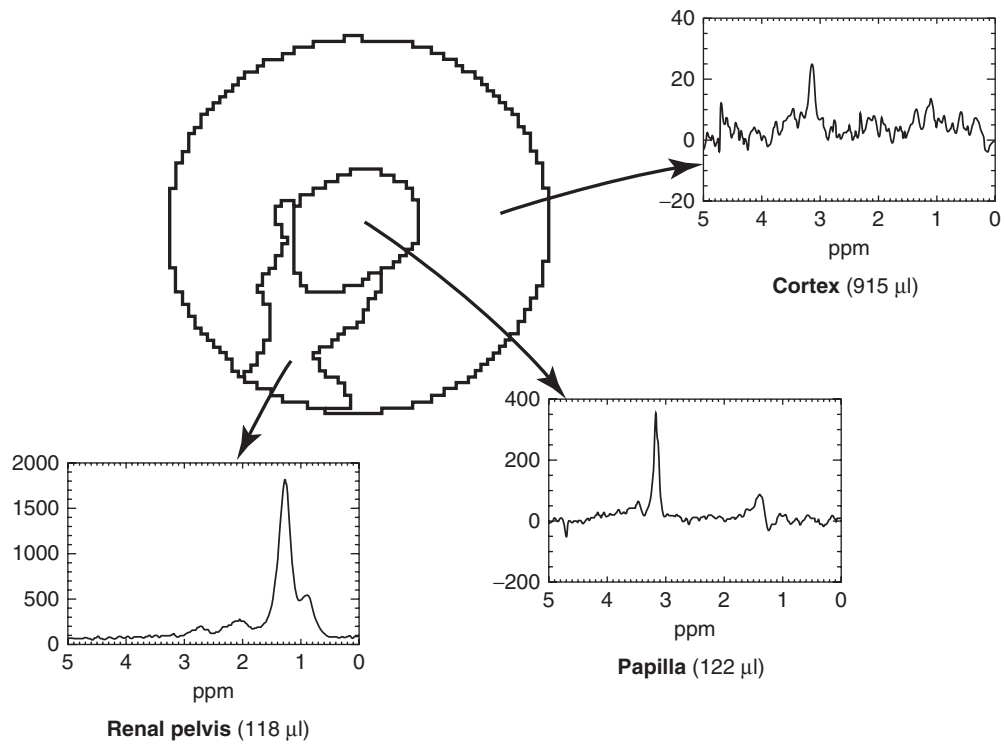

Figure 5. Localized ^{1}H-SLOOP spectra obtained in excised rabbit kidney in the experimental setup shown in Figure 3: the renal pelvis contains very strong lipid signals around 1.3 ppm, whereas the signal of trimethylamine (at 3.1 ppm) is detected in the cortex and the papilla. The high quality of localization is demonstrated by the good suppression of lipid signals in the latter two compartments. (Reprinted from *Journal of Magnetic Resonance* 94: 268–287; M. von Kienlin and R. Mejia: "Spectral Localization with Optimal Pointspread Function" (1991), with permission from Elsevier)

with γ the gyromagnetic ratio. The variable g_{mn} – just as in equation (6) in the case of linear phase-encoding – describes the signal contribution of the nth compartment, and is computed by integrating the local phase over the volume of the compartment:

$$g_{mn} = \int_{\text{compartment } n} e^{-\mathrm{i}\cdot\varphi_m(\vec{r})} \cdot \mathrm{d}^3\vec{r} \qquad (11)$$

The reconstruction process for obtaining localized spectra remains identical to the linear case: a pseudoinverse of the matrix **G** is computed, and the detected signals are combined with the complex reconstruction coefficients.

It has been demonstrated in phantom studies that SLOOP^{N} encoding with higher-order gradients can produce superior localization performance.[11] Coils that are positioned on the outside of the sample nevertheless have a limited ability to generate very steep local gradient fields. This approach therefore works best when only very few and large ROIs that lie relatively close to the magnet system's isocenter are to be selected within the sample volume.

Discussion and Conclusions

A few practical remarks are warranted. First, it should be pointed out that the above assessment is primarily theoretically based. Just as in spectroscopic imaging (CSI) – or any other MRS protocol for that matter – the spectral quality that can be achieved with SLOOP depends on the stability of the system and the experimental setup. Motion of the patient is as much a concern as for any other technique, as it will lead to subtraction errors that affect the accuracy of the result. All possible measures to minimize or compensate for residual patient motion are advised. As noted earlier, the ease of phasing of SLOOP spectra with zero- and first-order phase coefficients is a fair indicator of the quality of localization. Whenever the phase changes from one spectrum to the other or when spectra can only be looked at in 'magnitude' mode, there is likely to be some unaccounted for effect that should raise suspicion.

The optimization of individual phase-encoding gradients described in the section titled 'Spatial Response Function: The Optimal Experiment' involves some significant computational power to determine the best value in a high-dimensional space, which was not practical in the early SLOOP studies while a subject was waiting in the instrument. Experience has taught us that the optimal solution typically approximates to an acquisition-weighted bell-shaped distribution of the phase-encoding steps, just as has been proposed for conventional CSI.[12,13] A much faster optimization procedure thus can consist of fixing a constant distribution shape of the samples in k-space and optimizing only two (or three) parameters of nominal spatial resolution in each direction. Finally, for replicated measurements in setups with similar geometries – say a ^{1}H brain study or a ^{31}P cardiac exam – a single optimization may well provide the best solution for each measurement, which would avoid optimization for every individual case.

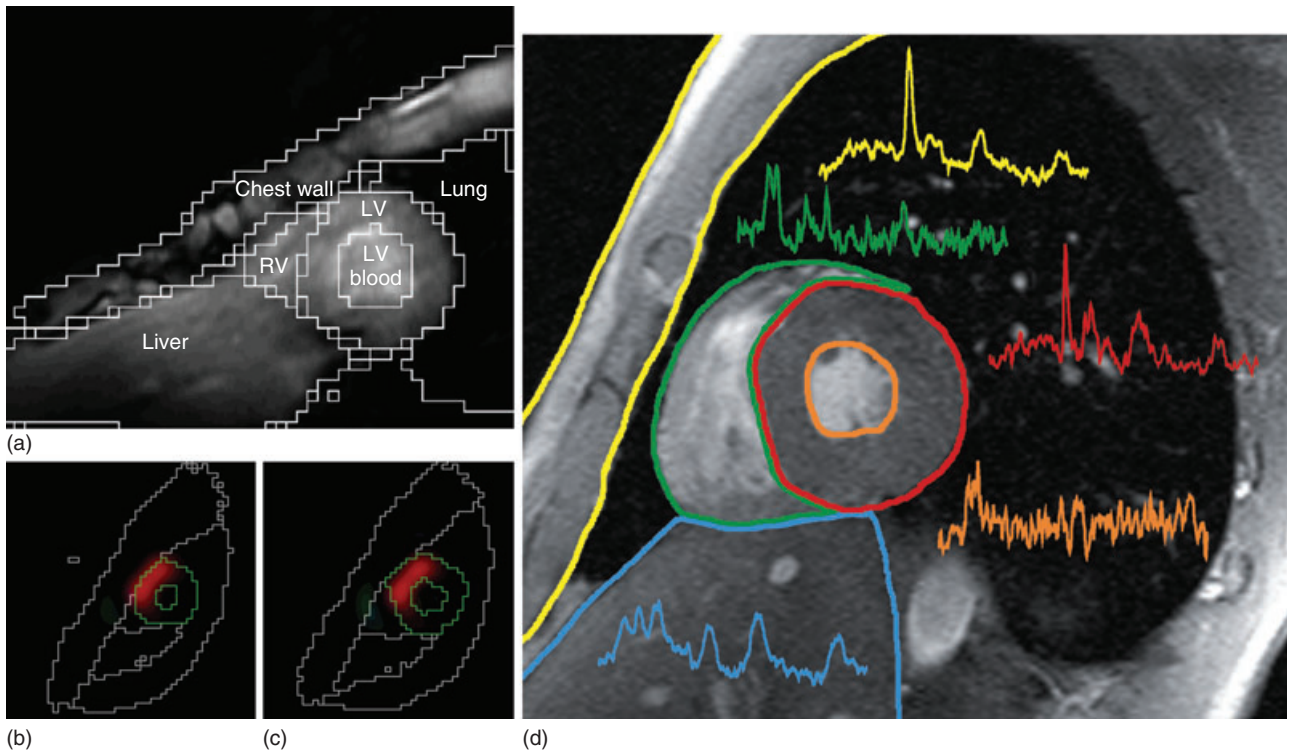

Figure 6. Application of SLOOP for ^{31}P-MRS in the human heart. (a) A priori information on the main anatomical compartments can be obtained on structural images. (b, c) Two planes of the SRF for the left ventricular wall. The decreasing amplitude toward the dorsal portion of the heart is due to the decreasing sensitivity of the surface coil, which has been accounted for in the model of this experimental setup. (d) Anatomically matched localized spectra, color coded for the various compartments chest wall (yellow), left ventricle (red), liver (blue), left ventricular blood (orange), and right ventricle (green). (Reprinted from *Journal of Magnetic Resonance* 134: 287–299; R. Löffler, R. Sauter, H. Kolem, A. Haase and M. von Kienlin: "Localized Spectroscopy from Anatomically Matched Compartments: Improved Sensitivity and Localization for Cardiac 31P MRS in Humans" (1998), with permission from Elsevier)

SLOOP has another very convenient property: localized spectra are by definition scaled to unit volume. Even if resonance signals have been detected in ROIs of largely different size, their spectra can be directly compared in terms of tissue concentration; the necessary spatial scaling is inherent in the reconstruction process.

In conclusion, SLOOP is an example of how a priori information can be successfully utilized to improve the accuracy of an experimental MRS assay. Ultimately, experimental success will depend on the conduct of the experiment including the data acquisition parameters: processing methods can only recover information present in the measured data, and cannot compensate for suboptimal experimental conditions. With the criteria presented for localization quality and SNR efficiency, and a recipe for obtaining an optimal set of experimental parameters, SLOOP may help advance another notch on the frontier of spatial resolution in MRS.

Acknowledgments

Robert S. Balaban at the National Institutes of Health in Bethesda started this work in 1990 by introducing me to the SLIM publication. I also thank my other main collaborators in this project at the NIH, in particular Raymond Mejia and Chrit Moonen. Over the years, this work was mostly pursued by many colleagues at the University of Würzburg, and I am very grateful, in particular, to Ralf Löffler, Wilfried Landschütz, Martin Meininger, Rolf Pohmann, Stefan Neubauer, and Meinrad Behr for their many insights and great work.

Biographical Sketch

Markus von Kienlin, b 1960; B. El. Engineer. 1982 Munich; PhD 1988 Grenoble. Was exposed to a 31P-MR spectrum of rat brain for the first time in 1984 and started to develop innovative methods for localized in vivo spectroscopy in animals and man. In pharmaceutical R&D at F. Hoffmann-La Roche Ltd., Basel, since 1999. Published over 100 manuscripts on MR procedures and their biological applications. Current scientific focus: brain function affected by disease and treatment.

Related Articles

Encoding and Decoding with Prior Knowledge: From SLIM to SPICE; CSI and SENSE CSI; Accelerated Spatially Encoded Spectroscopy of Arbitrarily Shaped Compartments Using Prior Knowledge and Linear Algebraic Modeling

References

1. D. I. Hoult, S. J. W. Busby, D. G. Gadian, G. K. Radda, R. E. Richards, and P. J. Seeley, *Nature*, 1974, **252**, 285.

2. D. I. Hoult and R. E. Richards, *J. Magn. Reson.*, 1976, **24**, 71.
3. X. P. Hu, D. N. Levin, P. C. Lauterbur, and T. Spraggins, *Magn. Reson. Med.*, 1988, **8**, 314.
4. R. Löffler, R. Sauter, H. Kolem, A. Haase, and M.von Kienlin, *J. Magn. Reson.*, 1998, **134**, 287.
5. M.von Kienlin and R. Mejia, *J. Magn. Reson.*, 1991, **94**, 268.
6. M. von Kienlin, K. Karino, K. Mori, P. Kador, D. Despres, J. Fleming and R. S. Balaban, *10th Annual Meeting of the Society for Magnetic Resonance in Medicine*, 1991.
7. M. Meininger, W. Landschutz, M. Beer, T. Seyfarth, V. Horn, T. Pabst, A. Haase, D. Hahn, S. Neubauer, and M.von Kienlin, *Magn. Reson. Med.*, 1999, **41**, 657.
8. H. Koestler, W. Landschuetz, S. Koeppe, T. Seyfarth, C. Lipke, J. Sandstede, M. Spindler, M.von Kienlin, D. Hahn, and M. Beer, *Magn. Reson. Med.*, 2006, **56**, 907.
9. M. Beer, T. Seyfarth, J. Sandstede, W. Landschuetz, C. Lipke, H. Koestler, M.von Kienlin, K. Harre, D. Hahn, and S. Neubauer, *J. Am. Coll. Cardiol.*, 2002, **40**, 1267.
10. W. Machann, F. Breunig, F. Weidemann, J. Sandstede, D. Hahn, H. Koestler, S. Neubauer, C. Wanner, and M. Beer, *Eur. J. Heart Fail.*, 2011, **13**, 278.
11. R. Pohmann, E. Rommel, and M.von Kienlin, *J. Magn. Reson.*, 1999, **141**, 197.
12. T. H. Mareci and H. R. Brooker, *J. Magn. Reson.*, 1991, **92**, 229.
13. R. Pohmann and M.von Kienlin, *Magn. Reson. Med.*, 2001, **45**, 817.

Further Reading

P. Bluemler, B. Bluemich, R. Botto, and E. Fukushima (eds), Spatially Resolved Magnetic Resonance, Wiley: Weinheim, 1998.
R. A.de Graaf, In Vivo NMR Spectroscopy, John Wiley & Sons Ltd: Chichester, 1998.