

CSI and SENSE CSI

Michael Schär¹, Bernhard Strasser² & Ulrike Dydak^{3,4}

¹ Johns Hopkins University, Baltimore, MD, USA

² Medical University of Vienna, Vienna, Austria

³ Purdue University, West Lafayette, IN, USA

⁴ Indiana University School of Medicine, Indianapolis, IN, USA

Chemical shift imaging (CSI) is a magnetic resonance spectroscopy (MRS) method of localizing spectra from multiple locations at the same time. For localization, the principle of phase encoding as used in MRI is applied. Because the applied phase-encoding gradients are strong and short, CSI localization does not suffer from image warping and chemical shift displacement artifacts. The main disadvantage of CSI is the long acquisition time. For example, acquiring the minimum of a single encoding step per phase encode leads to an acquisition time of more than 34 min for a two-dimensional acquisition with a 32×32 matrix and a typical 2 s repetition time. The actual spatial resolution can be described with a point spread function, which exhibits strong side lobes due to the limited extent of the encoding matrix and tissue heterogeneity, leading to contamination from neighboring voxels. Filtering reduces the side lobes and their associated contamination at the cost of resolution and signal-to-noise ratio (SNR) per unit volume. Acquisition weighting can be used to simulate a filter during acquisition and thus optimize SNR. Inhomogeneities in the static magnetic field cause shifts in the frequency domain of spectra from voxels at different locations. Those shifts can be corrected using a field map or the water signal from an additional CSI acquisition without water suppression. Water and lipid suppression are performed with the same approaches as in single volume acquisitions. As MRS is challenged by the low metabolite concentrations, phased-array coils are often used to improve the SNR compared to volume coils. Because of the nonuniform sensitivity profiles of each coil, the signal amplitudes and phases from different coils vary among voxels from different locations. Combining data from different channels can be done using the initial points of the FIDs from each voxel. An SNR-optimized coil combination can be achieved using sensitivity maps, and the signals in the different voxels scaled to display a homogeneous sensitivity distribution. Acquisitions with phased arrays can also be accelerated with parallel imaging methods, typically shortening the above-mentioned 34 min acquisition to 9 min or less.

Keywords: chemical shift imaging (CSI), spectroscopic imaging, phase encoding, multi-voxel spectroscopy, MRS, spectroscopy, parallel imaging, SENSE, GRAPPA

How to cite this article:

eMagRes, 2016, Vol 5: 1291 – 1306. DOI 10.1002/9780470034590.emrstm1454

Introduction

In magnetic resonance spectroscopy (MRS), the frequency information of the acquired signal is used to identify different chemical compounds. To achieve this, data are sampled over time (see *The Basics*). Unlike MRI techniques and with the exception of high-speed spectroscopic imaging (SI) (see *High-Speed Spatial-Spectral Encoding with PEPSI and Spiral MRSI*), a readout gradient is not applied during data acquisition. In the most basic pulse-and-acquire sequence, a free induction decay (FID) signal is recorded by a receive coil immediately after a single RF pulse is emitted by a transmit coil, without any gradients applied. In this experiment, the origin of the measured signal is defined by the combined spatial sensitivities of the transmit and receive coils. This is fast and simple, and is still used, for example, in dynamic muscle exercise studies of high-energy phosphate metabolism with ³¹P MRS.^{1–3} However, the localization provided has no sharp borders and is incapable of isolating signals from organs deeper in the body such as the brain or heart.

The so-called ‘single voxel’ or ‘single volume’ localization techniques were developed to address this need. These provide spatially selective excitation using RF (B_1^+) or static (B_0) field gradient pulses (see *Single-Voxel MR Spectroscopy; Localized MRS Employing Radiofrequency Field (B_1) Gradients*). Today, most single voxel localization techniques are based on B_0 gradient methods whose main advantages are the well-defined localization compared to the simple pulse-and-acquire method, and the fact that both B_0 and RF fields can be locally optimized for the narrowest line-widths and best water suppression. However, they also have disadvantages: (i) imperfect slice profiles can reduce the signal-to-noise ratio (SNR) and cause signal contamination from outside the selected volume; (ii) the selected voxels for different metabolites are spatially displaced due to their different chemical shift frequencies (see *Single-Voxel MR Spectroscopy*); and (iii) the MRS information is obtained from only a single voxel at a time.

Here, we present the multi-voxel localization method – chemical shift imaging (CSI), which overcomes the latter disadvantage. CSI localizes spectra from multiple locations simultaneously, enabling metabolic characterization of entire organs




Figure 1. CSI localization with phase encoding can be combined with any kind of selective or nonselective excitation method. Shown are (a) 3-D CSI with nonselective excitation, (b) 3-D CSI with volume-selective STEAM excitation, (c) 1-D CSI with a spin-echo-based 2-D column excitation, and (d) 2-D CSI with a slice-selective ISIS excitation. The phase-encoding gradients are shown in red

or regions of interest (e.g., whole brain slices) with the same SNR per unit time as single voxel methods. The simplest CSI approach employs the phase-encoding method from MRI to spatially localize the spectral information in an FID or spin echo with typical spatial resolutions ranging from 0.5 to 8 ml. Clinically, this enables the direct comparison of metabolic information from multiple regions of interest within a single scan. Challenges with this approach are (i) degradation of the spectral quality compared to single voxel methods because both B_0 and B_1^+ fields need to be optimized over a much larger volume; (ii) the long acquisition times required to phase-encode the entire volume; and (iii) the time required to process large multidimensional datasets. This article will explain the basic principles of CSI; the actual spatial resolution as described by the point spread function (PSF); how the PSF can be improved by either filtering or more SNR-efficient acquisition weighting; corrections for field inhomogeneities; water suppression; how to combine signals acquired with phased array detectors; and accelerated CSI acquisitions employing parallel imaging methods.

Basic CSI

Phase Encoding

CSI, also called SI, acquires spectroscopic data from a group of voxels (see **Chemical Shift Imaging**) using phase-encoding as in MRI (see **Image Formation Methods**). The main difference from MRI is that each data-sampling window in MRS is not used for spatial frequency encoding, but for encoding spectral frequencies and that phase-encoding is used to localize in 1-, 2-, or 3-dimensions (1-D, 2-D, or 3-D) to form a tissue

column, slice, or volume of voxels, respectively. The phase-encoding mechanism localizes the signal from all excited tissues, and can be combined with any type of signal excitation. For example, it may be desired to excite the entire object with a single non-selective excitation pulse, and in other cases a column, slice, or subvolume may be selectively excited using a single volume localization technique such as ‘point-resolved spectroscopy’ (PRESS),⁴ ‘stimulated echo acquisition mode’ (STEAM),⁵ or ‘image-selected in vivo spectroscopy’ (ISIS)⁶ (Figure 1).

Phase encoding requires spatially linear gradients. MRI scanners are equipped with three independent gradient coils that add linear variations in space to the otherwise homogeneous B_0 (which by convention defines the z -axis of the coordinate system). These gradient coils generate fields with linear gradients with amplitudes $G_x = dB_z/dx$, $G_y = dB_z/dy$, and $G_z = dB_z/dz$ in the B_0 field (see **Gradient Coil Systems**). The gradients establish a linear relationship between a spatial location (x, y, z) , the local magnetic field strength $B_z(x, y, z)$, and hence the local Larmor resonance frequency $\omega(x, y, z) = \gamma B_z(x, y, z)$, with γ , the gyromagnetic ratio, as shown in Figure 2. The phase-encoding gradient is a short pulse of one or more of these gradients applied immediately after a coherent signal has been excited with the RF pulse.

To illustrate the principle of phase encoding consider two samples in the scanner: sample A at the isocenter and sample B at some distance to its side (Figure 3a). In a first acquisition, both samples are excited and no phase-encoding gradient is applied. The spins of both samples precess at the same resonance frequency, so their signals, r^A and r^B , coherently add to produce $r_1 = r^A + r^B$ (Figure 3b). Now suppose that




Figure 2. The spatial dependence of the magnetic field and the local resonance frequency on an applied linear gradient: Shown are the static field B_0 (dotted), which is constant throughout, here along the spatial direction x ; and a linearly varying field from the superposition of the static field and an applied gradient G_x (solid). The local resonance frequency ω is given by the multiplication of the gyromagnetic ratio γ and the local magnetic field. With a linear gradient G_x applied, ω depends linearly on the location x

in a second acquisition both samples are again excited but a phase-encoding gradient is also applied (Figure 3b). During the phase-encoding gradient, the resonance frequency now differs between the two samples: sample A still resonates at the same frequency as in the first acquisition because the gradient does not add any field at the isocenter, but sample B resonates at a higher frequency because the phase-encoding gradient increases the local resonance frequency there. During the interval in which the phase-encoding gradient pulse is applied, the higher resonance frequency of B results in the accumulation of a net signal phase difference ϕ , compared to the signal from A. If the strength and duration of the phase-encoding gradient are selected to induce a phase difference of exactly 180° at B, then the signal of sample B will be opposite in sign to that of A, and the net signal from the second acquisition is $r_2 = r^A - r^B$. The signal amplitude of sample A, which reflects the number of spins at this location, can be determined by adding the signal amplitude (in reality a complex time-domain signal) of the two acquisitions ($r_1 + r_2 = 2 \times r^A$), and the signal

amplitude of sample B can be calculated by subtracting the signal amplitudes of the two acquisitions ($r_1 - r_2 = 2 \times r^B$).

Using the same principle, any number n of equally distributed locations along a spatial direction can be differentiated from n acquisitions in which a phase-encoding gradient pulse of different strength but equal duration is applied. However, with many samples a different strategy is needed to solve the equations for the signal at each location. Mathematically, the signal over time in an MRS experiment with 1-D spatial encoding (x), neglecting relaxation effects and suppressing the chemical shift information, is

$$s(t) = \int \rho(x) e^{i\phi(x,t)} dx \quad (1)$$

where $\rho(x)$ is the effective spin density along x , and ϕ is the accumulated phase. Using the counterclockwise positive sign convention,

$$\phi(x, t) = - \int_0^t \omega(x, t') dt' \quad (2)$$

where $\omega(x, t)$ is the local resonance frequency given by the sum of the Larmor frequency ω_0 , due to B_0 alone and an additional component induced by the temporally varying gradient field G_x :

$$\omega(x, t) = \omega_0 + \gamma x G_x(t) \quad (3)$$

In reality, other effects such as field inhomogeneities or eddy currents will add phase as well, but are neglected here.

After demodulation of the carrier frequency ω_0 , the phase accumulated due to the applied phase-encoding gradient of strength G_{PE} and duration τ is:

$$\phi_{PE}(x, G_{PE}) = - \int_0^\tau \gamma x G_{PE} dt' = -\gamma x G_{PE} \tau \quad (4)$$

Borrowing the k -space notation from MRI to define a spatial frequency

$$k(G_{PE}) = \frac{\gamma}{2\pi} G_{PE} \tau \quad (5)$$




Figure 3. Using 1-D CSI to distinguish two phantoms. (a) Phantom A is at the isocenter and phantom B, with less content, is off center. A phase-encoding gradient G_{PE} generates a linearly increasing field along direction x . (b) Two scans are performed, scan 1 without and scan 2 with G_{PE} applied. In scan 1 signals from both phantoms are in phase and add up to the combined signal of scan 1 (r_1). In scan 2, G_{PE} does not affect phantom A (because the gradient is zero at the iso-center) but its duration and strength are set to induce 180° phase in the signal of phantom B (r^B). Hence, r^B is subtracted from that of phantom A (r^A) in the combined signal of scan 2 (r_2). Signal from phantom A can then be determined by adding r_1 and r_2 , and signal from phantom B by subtracting them

the signal can be written as

$$s(k) = \int \rho(x) e^{-2\pi i k x} dx \quad (6)$$

which is the FT of the spin density. Therefore, the spin density of the sample can be calculated using the inverse FT of the signal:

$$\rho(x) = \int s(k) e^{2\pi i k x} dk \quad (7)$$

Because we are sampling k -space discretely, the discrete FT (DFT) must be used and the reconstructed spin density $\rho'(x)$ (with a prime to distinguish it from the continuous spin density $\rho(x)$) is given by

$$\rho'(x) = \Delta k \sum_{p=-(N/2)}^{(N/2)-1} s(p \Delta k) e^{2\pi i p \Delta k x} \quad (8)$$

where $s(p \Delta k)$ is the measured signal at the k -space location denoted by integer p , sampled at N different locations. Discrete sampling means the continuous k -space is multiplied by a comb function $u(k)$, and the reconstructed $\rho'(x)$ can be described by the convolution of $\rho(x)$ with $U(x)$, which is the inverse FT of $u(k)$ and a comb function as well. The result of a convolution with a comb function is a series of copies of the original spin density. To prevent aliasing of these copies, the Nyquist sampling criterion for the field-of-view (FOV),

$$\frac{1}{\Delta k} \geq \text{FOV} \quad (9)$$

has to be satisfied. The 'nominal' spatial resolution or Fourier voxel size, Δx , is also derived from the FOV and the maximum spatial frequency $k_{\max} = N/2 \Delta k$:

$$\Delta x = \frac{\text{FOV}}{N} = \frac{2}{k_{\max}} \quad (10)$$

When phase-encoding is applied in two or three spatial directions, a 2-D or 3-D FT is used to reconstruct the spatial information. For CSI, an FID or spin-echo with m time-point samples is acquired after each phase-encoding gradient pulse has ended. During this acquisition the different frequency components of the MRS signal evolve according to the chemical shifts of the moieties present (see **The Basics**), forming m k -space vectors: one with the first sample from all FIDs or echoes, one with the second sample from all the signals, etc. After a 1-D, 2-D, or 3-D FT in all of the spatial dimensions for all m k -space vectors, an FID or echo time (TE)-domain signal is available corresponding to each spatial location. These can be further processed to form a spectrum at each location. Instead of m spatial k -space vectors, it is common practice to treat the spectroscopic dimension simply as another dimension of k -space, k_p , sampled in the time domain with m data points.

A major advantage of phase-encoding over frequency-encoding, in which location is encoded with gradients applied during the signal readout, is the lack of chemical shift and off-resonance (B_0) warping artifacts. With spin warp phase-encoding, introduced by Edelstein *et al.* (see **Spin Warp Data Acquisition**),⁷ off-resonance warping is eliminated because

the phase-encoding gradients are much stronger than local B_0 gradients. Chemical shift artifacts are negligible because the frequency differences due to chemical shift are only ppm of the phase encoding gradient pulse strength. The absence of chemical shift artifacts is a key advantage of CSI, as it allows correct localization of wide ranges of chemical shift dispersions.

In practice, when CSI is used in the brain, it is often combined with single volume localization schemes employing selective excitation (see **Single-Voxel MR Spectroscopy**) to spatially confine MRS signals to a desired volume of interest (VOI). It is important to remember that, like readout gradients, frequency-selective RF pulses are subject to chemical shift displacement artifacts wherein signals from each metabolite derive from slightly shifted volumes (Figure 4). In contrast to single voxel MRS, the VOIs in CSI are usually larger, which increases the absolute amount of the chemical shift displacements. The chemical shift artifact can result in missing peaks from metabolite moieties at the edge of the spectrum and the border of the VOI. In contrast, the localization of each voxel in the CSI phase-encoding dimensions will have no such artifacts. In the center of the VOI all metabolite signals are correctly localized (PSF effects notwithstanding; see section titled 'Point Spread Function (PSF)'). The chemical shift displacement artifact cannot be corrected by post-processing, but a single




Figure 4. The combination of CSI localization with single volume localization such as PRESS or STEAM. A large volume of interest may cause chemical shift displacements, i.e., shifts of the excitation volumes for each metabolite with respect to each other, which are much larger than those occurring with single voxel MRS. On most scanners the planned VOI is on-resonance and thus displayed correctly for NAA (N-acetyl aspartate; blue, central VOI in this figure), whereas the excitation volumes for choline and total creatine are shifted toward one side of NAA, while the lactate and lipid VOIs are shifted to the other side according to their chemical shifts relative to NAA. However, it is important to note that the metabolite assignment of each CSI voxel achieved through phase-encoding (white squares) is *not* affected by such shifts

volume localization approach that is less prone to the artifact is discussed in section titled 'Water and Lipid Suppression Methods'.

Signal-to-noise Ratio (SNR) in CSI

Like single voxel MRS, a limiting factor in CSI is the low SNR due to the low metabolite concentrations and sensitivity of MRS methods in general. As in single voxel MRS, the SNR of a CSI spectrum is mainly governed by the effective voxel volume and the number of signal averages. In a CSI protocol, the number of averages for computing SNR is the number of averages per phase-encoding step times the number of different phase-encoding steps since the signal from the whole VOI is acquired in each. Thus, a CSI matrix of 16×16 will result in 256 averages, while a 32×32 matrix contributes 1024 averages toward the SNR of each spectrum. In fact, ignoring the effects of shimming over the CSI VOI as compared to the much smaller single voxel, a CSI voxel of effective size v , acquired with n averages equal to the number of phase-encoding steps times the number of averages per phase-encode, has the same SNR as a single voxel spectrum acquired from a volume v with n averages.

Point Spread Function (PSF)

A closer look at the spatial resolution in CSI shows that the actual voxel size can substantially deviate from the nominal Δx of equation (10) due to the Fourier encoding process. Only a limited amount of k -space is sampled discretely in CSI, which affects spatial resolution. The PSF is generally used to describe these effects. The PSF describes the response of an imaging system to a point source. In CSI, the PSF can be determined from the FT of the sampling function, and is exemplified for a 1-D acquisition with 16 phase-encoding steps in Figure 5a. In CSI the number of phase-encoding steps is generally small compared to MRI because of the relatively long repetition times (TRs) used in MRS, the need for larger voxel sizes due to the low

SNR, and the fact that only one spatial k -space point is sampled per acquisition.

The PSF shown in Figure 5b is determined by taking the FFT of the sampling function. To demonstrate the repetitive nature of the PSF outside the encoded CSI FOV, zeroes were added in between the acquired phase-encoding steps in the sampling function. If signal is excited outside the FOV, the repeating main lobes of the PSF (due to finite discrete sampling) will lead to aliasing artifacts. Adding zeroes at higher k -space positions in the sampling function interpolates the PSF to show its response at a higher resolution than the nominal voxel size Δx , revealing strong side lobes that can lead to significant contamination between neighboring voxels, also called 'voxel bleed'. The third observation is that the full width at half maximum (FWHM) of the PSF is 21% larger than the nominal Δx . The side lobes of the PSF can be reduced in post-processing by applying apodization functions in k -space, albeit at a cost to spatial resolution and SNR per unit time. Figure 6 shows the sampling function and its corresponding PSF for a 1-D CSI acquisition with 16 phase-encoding steps with and without filtering using a Hanning function. Filtering significantly reduces the side lobes and therefore the contamination from neighboring voxels. However, the FWHM of the filtered PSF is about twice the nominal Δx .

Acquisition-Weighted CSI

Instead of applying a post-processing filter, one can spend more time acquiring data at the center of k -space and less time at its periphery. This is called 'acquisition weighting'. Filtering, or weighting, k -space data in post-processing is less efficient in terms of SNR per unit time, compared to acquisition weighting.^{8,9} 2-D proton (^1H) CSI is usually acquired with matrix sizes of 16×16 to 32×32 , leading to 256–1024 acquisitions, respectively. With a TR of 2 s, the total acquisition times range from 8.5 to 34 min when performed with one average per phase encode. A simple acquisition weighting often employed for 2-D CSI is a circular shutter wherein the k -space corners




Figure 5. (a) Sampling function for a 1-D CSI acquisition with 16 phase-encoding steps. The sampling function is set to 1 for the 16 acquisitions, and 0 for both higher k -space positions and in between the 16 encoding steps. (b) The PSF (black), which is the FFT of the sampling function, shows where the signal from a point source is reconstructed. The PSF is normalized to 1 at the center of the target voxel, and is zero at the center of the nominal voxel locations, but has significant side lobes next to the target voxel, which leads to significant signal contamination between neighboring voxels when the point source is off center. Because k -space is sampled discretely and not continuously, the PSF repeats itself outside the FOV, which generates aliasing artifacts if there is any signal present outside the encoded FOV. The nominal voxel location is shown in gray. The FWHM of the PSF is 21% wider than the nominal voxel size, Δx .




Figure 6. (a) Sampling functions with and without Hanning filtering for a 1-D CSI acquisition with 16 phase-encoding steps. (b) The normalized PSF from the sampling functions in (a) together with the nominal voxel location. Filtering significantly reduces the side lobes at the expense of an increase in FWHM




Figure 7. (a) Sampling functions for 2-D CSI acquisitions with 16 phase-encoding steps in both directions. The circles indicate the acquired k-space locations. Those without gray filling outside the large black circle are omitted when a circular acquisition weighting shutter is applied. (b) The PSF without the shutter shows side lobes along the main grid directions but reduced lobes along the diagonals. (c) With the circular shutter, the PSF is circularly symmetric

are not acquired (Figure 7a). Figure 7b and c shows the PSF of 2-D CSI acquisitions with 16 phase encodes without and with a circular acquisition-weighting shutter. Without the circular shutter (Figure 7b), side lobes are strong only along the two grid directions, while they are much reduced along the diagonal directions. This is because the k -space corners are further away from the center, leading to more data points and higher spatial frequencies being sampled along the diagonal directions. Thus, the resolution and the side lobes of the PSF are improved along the diagonals. Applying a circular shutter (Figure 7c) makes the PSF circularly symmetric with the same resolution in all directions. The same weighting can be performed for 3-D CSI as well, using a spherical or ellipsoidal shutter. Not acquiring the corners of k -space leads to time-savings of roughly 21% and 48% for the 2-D and 3-D CSI experiment employing a single average per phase-encode, respectively.

In non- ^1H CSI, such as ^{31}P and carbon (^{13}C) CSI, the voxel size Δx is usually chosen larger than in ^1H CSI because of the

lower sensitivity. Larger voxels reduce the number of phase-encoding steps required and lead to shorter acquisition times. To further improve SNR, signal averaging is often applied. With multiple signal averages per phase-encode, the number of averages acquired per k -space location can be varied. In the 'accumulation weighted' acquisition scheme, more averages are acquired at the center and fewer at the edge of the k -space (Figure 8).^{8,9} The PSF from accumulation weighting is similar to that of uniform sampling with filtering applied in post-processing, but without loss of SNR.

A third acquisition weighting method is 'density weighting'.¹⁰ Density weighting does not require the acquisition of multiple averages and is therefore more flexible. Density-weighting varies the distance between neighboring sampling points to approximate the desired weighting function (Figure 8a). No data weighting is required in post-processing and SNR loss is avoided as in accumulation weighting. In addition, because sampling is more dense at the center of




Figure 8. (a) Sampling functions for 1-D CSI acquisitions with 45 acquisitions. Depicted are uniform sampling with three signal averages (light gray); accumulation weighting with varying number of averages approximating a Hanning function (dark gray); and density weighting with one average each distributed nonuniformly in k -space at locations that approximate the Hanning function (black). The corresponding PSFs are shown in (b). Accumulation weighting reduces side lobes similar to filtering in post-processing (Figure 6) but without the loss in SNR. Density weighting reduces the side lobes and additionally extends the FOV as k -space is sampled more densely at the center of k -space, but requires a more complex reconstruction as data have to be re-gridded first

k -space, the effective FOV is extended compared to uniform or accumulation-weighted sampling (Figure 8b). However, because the data are not acquired on a regular Cartesian grid in k -space, they need to be projected onto a regular Cartesian grid to perform FFT and a roll-off correction afterwards.¹¹

B_0 Correction Methods

A drawback of CSI compared to single voxel MRS is the increase in B_0 field inhomogeneity resulting from the data being collected from a much larger volume. In CSI, B_0 shimming is performed to decrease B_0 inhomogeneities across the full CSI VOI, while in single voxel MRS, strong local shimming can potentially yield sharper spectral peaks, as noted earlier. Reducing the voxel size in CSI increases the local homogeneity within each voxel and thus the inhomogeneously broadened spin-spin decay constant (T_2^*). Therefore, the problem of local shimming diminishes as the spatial resolution approaches that of MRI. In fact, the increased local homogeneity per CSI voxel may in part compensate for the signal loss due to the volume decrease. Li *et al.* have shown that in the brain at 4 T, the volume of a cubic voxel decreases by the third power of the voxel dimension, while the SNR reduces only quadratically due to T_2^* increases.¹²

In CSI, the B_0 field inhomogeneity across the FOV leads to shifts in resonances from voxel to voxel. This is not a concern if the time- or frequency-domain fitting can accommodate the variations in metabolite frequency. Such post-processing is often more robust if the spectra from different voxels are first frequency-aligned on the basis of a B_0 field map or a strong signal from either residual water or *N*-acetyl aspartate (NAA). Robust frequency alignment typically utilizes an additional CSI dataset acquired without water suppression. The water signal at each voxel location can not only be used for frequency alignment, but also for automatic correction of phase,

including eddy current-induced phase modulations during the acquisition window that lead to line-shape distortions. Since B_0 -induced frequency shifts, spatially varying phase, and line-shape distortions are the same for all spectral peaks in one voxel, the phase of the water signal in the non-suppressed acquisition can be subtracted from the water-suppressed data to automatically perform the frequency alignment, a zero-order phase correction, and eddy-current correction.^{13,14}

Water and Lipid Suppression Methods

In ^1H MRS the water signal is usually suppressed because side bands and baseline distortions caused by the two to four orders-of-magnitude stronger water signal distort the metabolite signals (see *The Basics; Single-Voxel MR Spectroscopy*). The most commonly used methods for water suppression apply chemically selective irradiation prior to MRS signal excitation to minimize the longitudinal magnetization of water at the time of excitation. ‘Chemical shift-selective’ (CHESS) RF pulses tip the longitudinal magnetization of water into the transverse plane, where it gets dephased by a gradient crusher pulse while the magnetization at other frequencies remains unaffected.^{15,16} Because of spin-lattice (T_1) relaxation, the applied flip angle should be larger than 90° so that a negative longitudinal magnetization can relax toward the zero-crossing during the time gap between the water suppression and signal excitation. Therefore, optimal suppression depends on the T_1 of water, the applied flip angle (which depends linearly on the local B_1^+ field), and the timing between the sequence elements.

Spatial variations in the water T_1 and B_1^+ inhomogeneities lead to varying suppression quality. To improve suppression, multiple CHESS pulses are often applied. The ‘water suppression enhanced through T_1 effects’ (WET) method applies three or four CHESS pulses with different flip angles.¹⁷ The flip angles are optimized on the basis of a Bloch equation




Figure 9. Pulse sequence diagrams of a PRESS localized acquisition with either (a) WET or (b) VAPOR water suppression. OVS sequence elements (gray) can be combined with both; however, in WET the water suppression CHESS pulses have to be applied earlier to accommodate the OVS elements in between WET and spatial localization. This may require re-optimization of the applied WET flip angles. With VAPOR, multiple OVS elements can be interleaved with the VAPOR CHESS pulses

analysis for ranges of T_1 relaxation times and B_1^+ amplitudes. WET performs robustly for single voxel MRS at 1.5 T. In CSI, however, the VOI is much larger compared to single voxel MRS, which typically results in a larger range of B_1^+ values depending on the uniformity of the transmit coil. B_1^+ inhomogeneities tend to increase with B_0 especially at 3 T and higher. The ‘variable pulse power and optimized relaxation delays’ (VAPOR) method applies six to eight CHESS pulses, which provide robust water suppression despite large B_1^+

inhomogeneities.¹⁸ A key difference between VAPOR and WET is the time required for their application: WET lasts about 100–200 ms, whereas VAPOR takes around 600–800 ms, which can be limiting in multislice CSI acquisitions. Sequence diagrams for WET and VAPOR are given in Figure 9.

In addition, even if the VOI does not contain large signals from lipids, surrounding tissues often do contain intense lipid resonances, e.g., lipids in the skull and scalp around the brain. Because these signals are much stronger than the metabolite




Figure 10. Chemical shift displacement in volume pre-localization versus OVS. To suppress subcutaneous lipid, (a) only tissue in the green box is excited using single volume localization, or (b) eight OVS slabs are placed on the skull around the brain while the tissue is excited in the entire slice. Resulting NAA, creatine, and choline metabolite maps are shown for each method. In (a), using PRESS localization at 3 T for large volumes of interest results in slightly shifted volumes of excitation for (total) creatine and choline signals compared to NAA due to chemical shift displacement in single volume localization. OVS (b) targets the lipid signal and hence does not suffer from chemical shift displacement, and allows more arbitrarily shaped volume selection




Figure 11. Planning options for a 2-D CSI brain acquisition with an 18×18 matrix (top row) and corresponding NAA maps (bottom row). (a) If only phase encoding together with slice selection is applied in low- and medium-resolution ^1H CSI, the NAA signals are overwhelmed by subcutaneous lipid, also resulting in signal artifacts inside the brain due to the large PSF side lobes (bottom). (b) To suppress signal from subcutaneous lipids, eight OVS slabs are placed on the skull around the brain (blue bars, top). The performance of OVS slabs with different suppression pulses may vary, and placing OVS slabs is sometimes tricky, leaving some lipid artifacts in the data (bright peripheral areas, bottom). In (c) only the tissue within the green VOI is excited using PRESS. Note that PRESS is subject to chemical shift displacement artifacts at 3 T. In this case, the slightly shifted VOI excited some lipid signals on the left side where the displaced lipid VOI meets the scalp (bottom). In (d) PRESS is adjusted to localize a larger region but is combined with OVS slabs

signals, even a little voxel bleed (see section titled ‘Point Spread Function (PSF)’ may contaminate the VOI. Therefore, CSI is usually also performed with lipid suppression. This can be achieved with CHESS pulses tuned to the lipid frequency, or more commonly, with either volume pre-localization (see section titled ‘Phase Encoding’, Figure 4) and/or with ‘outer volume suppression’ (OVS). In pre-localization, the single volume localization (see *Single-Voxel MR Spectroscopy*) of the VOI is positioned to avoid exciting the surrounding lipids (Figure 10a) so that they do not contribute to the measured signal. Note that when localizing with slice-selective RF pulses, imperfections in the slice profile will reduce the signal at the edge of the VOI, thereby blurring the borders of the metabolite maps. Moreover, the creatine and choline volumes are shifted compared to the on-resonance NAA map due to the chemical shift displacement (Figure 10a).

OVS suppresses all signals within a volume by applying highly selective slice-selective RF pulses followed by crusher gradients to dephase the excited signal. Multiple OVS slices can be used to carve out arbitrarily defined polygon-shaped volumes. Figure 10b shows how the lipid surrounding the brain can be suppressed with eight OVS pulses. Since the OVS pulses specifically target the lipid frequency, the localized suppression is not affected by chemical shift displacement. If OVS is applied with slice-selective excitation instead of volume pre-localization, the resulting metabolite images do not suffer from chemical shift displacement artifacts. NAA and creatine metabolite maps from acquisitions with both

slice-selective excitation and OVS exhibit an elliptically shaped VOI (Figure 10b).

Figure 11 demonstrates some practical considerations when planning a ^1H MRS CSI acquisition in the brain. If neither volume pre-localization nor OVS are used, the PSF side lobes of the large subcutaneous lipid signal will overshadow most other metabolite signals within the brain in low-to-medium resolution CSI (i.e., when using voxel sizes larger than about $5 \times 5 \times 5 \text{ mm}^3$). In Figure 11b, eight OVS bands are placed to locally suppress subcutaneous lipids. Depending on the suppression pulses used, the performance of OVS may vary, as this example shows. Furthermore, placing of OVS bands may be tricky especially for 3-D CSI, often leaving some residual lipid artifacts in the data. Figure 11c shows the planning and the NAA map when using pre-localization with PRESS at 3 T. The VOI needs to be carefully planned to avoid the skull and air cavities such as the sinus, eyeballs, etc. (which are prone to susceptibility artifacts). Thus, the VOI is usually limited to more central parts of the brain rather than cortical areas. Note that even with the on-resonance NAA map centered within the planned VOI, the chemical shift displacement of the lipid VOI results in some lipid artifacts in the corners of this simple integration-generated metabolite map. As long as this unwanted lipid signal is not overwhelming, creating metabolite maps with spectral fitting instead of integration will generally take care of such artifacts. At clinical field strengths of 1.5 and 3 T, both approaches, volume pre-localization and OVS, are often combined as shown in Figure 11d. This allows a more

generous planning of the VOI that may include some lipid areas, as long as they are covered by OVS.

An advantage of VAPOR water suppression is that it allows the interleaving of multiple OVS modules with its CHESSE elements. Zhu *et al.* extended the VAPOR approach by using dual-band frequency-modulated CHESSE pulses that saturate both the water frequency and the methyl and methylene resonances (0.8–1.4 ppm) of lipid.¹⁹

CSI with Phased Arrays

In recent years, the use of phased arrays (see *Receiver Loop Arrays*) has become routine in MRI due to the improvements in SNR and the possibility of performing parallel imaging (see *Spatial Encoding Using Multiple rf Coils: SMASH Imaging and Parallel MRI*) that they offer. For CSI, the SNR improvement is especially important in metabolic studies because of its intrinsically low SNR per unit time. The challenge with data from phased arrays is that the different coils have spatially varying sensitivities both in magnitude and phase, which makes the combination of the data nontrivial. To combine data from different coils with maximum SNR, complex weights w_n have to be estimated at each voxel location. The combined signal s assembled from a linear combination of signals, s_n , from each individual coil channel is given by

$$s(\mathbf{r}, t) = \lambda(\mathbf{r}) \sum_{n=1}^N s_n(\mathbf{r}, t) w_n(\mathbf{r}) \quad (11)$$

where $\mathbf{r}$ is the spatial position, t are the FID time points, λ is a scaling factor, and N is the total number of coil channels. Here, the noise correlation between channels is neglected for simplicity.

According to Roemer *et al.*, the SNR is maximized by choosing the weights based on sensitivity maps from the coils.²⁰ In MRI, a simple sum-of-the-squares (SoS) reconstruction can be used instead of sensitivity-based optimization at a moderate SNR loss of about 10%, as long as the SNR is above a certain threshold. However, SoS results in magnitude data only, which is unsatisfactory for CSI, as magnitude spectra have a much broader linewidth. Therefore new coil data combination methods have been developed for CSI. These can be classified as either intrinsic or extrinsic. Intrinsic methods use the spectroscopy signal s_n itself to derive the weights w_n , whereas extrinsic methods use additional external data.

Intrinsic Methods

The basic intrinsic CSI coil combination technique determines $w_n(\mathbf{r})$ from the complex conjugate (denoted by $*$) of the first time point of the CSI FID of each coil, i.e., $w_n(\mathbf{r}) = s_n^*(\mathbf{r}, t = 0)$.²¹ The advantage of this method lies in its simplicity and the fact that no additional data are needed. Reported problems arise from the degraded SNR of the water peak as it is partially suppressed, phase problems from distorted water peaks, lipid contamination and the inherent phase difference between water and the metabolites of interest.^{22,23} Those disadvantages are overcome when applying the same method to non-water-suppressed CSI. However, the lack of water

suppression limits its application to spin-echo acquisitions with long TEs, since strong water sideband artifacts are often present at short TE without water suppression.²²

In another intrinsic method, data are weighted with the maximum of the magnitude spectra and are phased before signal combination by minimizing the difference between the absorption and the magnitude spectrum from a prominent peak.²⁴ Although this method may lead to good results for high SNR spectra, the phasing can be problematic in low SNR spectra.

The challenge with all intrinsic methods is the determination of the scaling factors, λ , which define the relative signal strengths in voxels close to the coils, as compared to those further away. This basically limits the application of intrinsic methods to metabolic ratio maps, where location-dependent signal weighting normalizes out in the ratio, unless absolute quantification is performed (see *Measuring Metabolite Concentrations I: ¹H MRS*).

Extrinsic Methods

The most common extrinsic coil combination method uses measured sensitivity maps. They require measurements from two images, one with the phased array and one with a homogeneous reference (volume or body) coil. Sensitivity maps are calculated by dividing the image of each channel by the image of the reference coil in a complex manner. This method achieves SNR-optimized coil combination if the noise correlation matrix is taken into account.²⁰ It is very efficient for ¹H MRS because the time for acquiring the sensitivity maps is typically negligible compared to the CSI acquisition.²⁵ If the reference image is acquired with a volume coil with a homogenous sensitivity over the entire object, the relative spatial scaling is also inherently taken care of, which is a significant advantage over all of the intrinsic methods. However, homogenous reference coils are not always available, e.g., at field strengths of 7 T or above, or when using multichannel transmit coils, or in multinuclei spectroscopy such as in ³¹P CSI.

If absolute quantification is performed using water as the internal concentration reference, an additional CSI data set must be acquired without water suppression.^{26,27} The coil combination weights can be computed from these data by determining the phase and amplitude of the water resonance in the spectral domain. As well as providing all the information needed for extrinsically combining the phased-array MRS signals, as discussed above, the additional non-water suppressed acquisitions allow automatic frequency alignment, zero-order phase and line-shape corrections, albeit at a cost to measurement time.

Parallel Imaging in CSI

Parallel imaging in MRI relies on the simultaneous detection of signal with multiple receiver coils (so called phased array coils). Each coil element has its own spatial sensitivity pattern, representing an intrinsic spatial encoding pattern that can be further exploited.²⁸ Uniformly omitting phase-encoding steps reduces the scan time, but spatial aliasing occurs as a result of the reduced FOV (Figure 12A–C). The individual




Figure 12. CSI with SENSE, phased-array detectors, and GRAPPA. (A) k -Space sampling pattern showing a fully sampled 16×16 k -space indicated by circles. With parallel imaging using an acceleration factor of 2 in each direction, only a quarter of the k -space locations (gray-filled) are acquired, resulting in a fourfold acceleration of the acquisition. The effect of the 2×2 -fold undersampling of the image at locations a–d in (B) leads to an aliased image with half the FOV as shown in (C). (D) Spectra from each location. The spectrum acquired from point e in (C) is aliased by a superposition of signals from the locations indicated by the gray dots in (C). In (E) SENSE uses the low-resolution sensitivity maps acquired with both the phased array and a homogenous volume coil to separate the aliased voxels into the original full-sized FOV voxels. (F) GRAPPA uses additional data acquired at the center of k -space called ‘Auto-Calibration Signals’ (ACS), shown in the undersampled CSI k -space. However, CSI often uses an additional calibration image acquired for each coil n . In the fully sampled ACS k -space, weights W are calibrated to determine the signal at the target location (white circle) from a linear combination of all neighboring source locations within the GRAPPA kernel (white box). In a second step, the weights are used to estimate the signal at an unmeasured target location (fat black circle) based on acquired source locations within the GRAPPA kernel (gray box)

sensitivities of the different receiver coils can then be exploited to reconstruct images that are free of aliasing artifacts. Because CSI uses phase encoding in all spatial directions, it can be accelerated in all of them, whereas in Cartesian imaging the readout direction cannot be accelerated by these means.

Two parallel imaging methods are commonly used in MRI. In the sensitivity encoding (SENSE) method, aliased images are unfolded in the image domain, based on the low resolution sensitivity maps introduced for phased array reconstruction (Figure 12E).²⁹ In the generalized auto-calibrating partially parallel acquisition (GRAPPA) method, missing data points are estimated in k -space for each coil based on the so-called ‘auto-calibration signal’ (ACS) data instead of sensitivity maps (Figure 12F).³⁰

SENSE

SENSE with a phased-array is a powerful tool for fast CSI.³¹ For a 2-D SENSE-CSI acquisition, the FOV is reduced by a factor R_x

in the x direction and a factor R_y in y direction. In this manner, only a fraction of the full k -space points is sampled, leading to a scan time reduction by a factor $R = R_x \times R_y$. Thus, if the full FOV is to be resolved by $n_x \times n_y$ spectra, only $n_x/R_x \times n_y/R_y$ individual signals need be acquired. As an example, the black circles in Figure 12A show the sampled k -space locations of a 16×16 CSI acquisition. By accelerating this acquisition with SENSE factors of 2 in both directions, only the k -space locations with gray filled circles need be acquired, resulting in a $2 \times 2 = 4$ -fold faster scan. As discussed in the section titled ‘Point Spread Function (PSF)’ (Figure 6), uniformly reducing the sampling density in k -space reduces the FOV in the image domain, which leads to aliasing if the FOV is smaller than the object within the coil’s sensitive volume. Figure 12B shows a fully sampled image, and Figure 12C is the image with half of the FOV in both directions showing the corresponding aliasing artifacts. The signal in the voxel (e) with the black dot is the superposition of the signal at its original location and the three voxels marked with

gray dots. Figure 12D shows the corresponding spectrum of the aliased voxel (e), the original correct spectrum (a), and the spectra from the aliasing locations (b–d).

The SENSE reconstruction exploits the fact that each signal contribution is weighted according to the local sensitivity of the respective coil. 2-D SENSE-CSI, with two spatial and one frequency dimension, is similar to 2-D SENSE MRI.³² The spectral encoding of the frequency dimension can be treated as an additional dimension without undersampling just like the spatial readout direction in SENSE MRI. After FFT in all dimensions, an aliased image needs to be unfolded for every frequency step in the spectral direction.

Consider one voxel in the reduced FOV (e in Figure 12C), and the corresponding set of voxels in the full FOV (a–d in Figure 12B). Using the notation of Pruessmann *et al.*, S denotes the sensitivity matrix, containing the complex spatial sensitivity of each coil for each superimposed voxel position.²⁹ Then the unfolding matrix U is given by

$$U = (S^H \Psi^{-1} S)^{-1} S^H \Psi^{-1} \quad (12)$$

where the superscript H denotes the transposed complex conjugate (adjoint operator), and Ψ is the receiver noise covariance matrix. The matrix Ψ is determined experimentally in a pre-scan, using

$$\Psi_{\alpha\alpha'} = \overline{\eta_\alpha \eta_{\alpha'}^*} \quad (13)$$

where η_α denotes noise of the α th receiver channel, the bar denotes time averaging, and the asterisk denotes the complex conjugate. For each frequency step λ , signal unfolding is achieved by

$$v_\lambda = U a_\lambda \quad (14)$$

where the vector a_λ lists the complex image values of the chosen voxel in the aliased images obtained by each coil, and the resulting vector v_λ lists the unfolded voxel values in the full FOV. This procedure is repeated for each voxel in the reduced FOV and for each λ in the spectral direction to obtain a full SI dataset. As a result, each spectrum gets unfolded and only shows metabolite peaks actually contained in the corresponding full-FOV voxel.

The SENSE reconstruction is performed only for voxels within the object border because the coil sensitivity outside the object cannot be measured, and because reducing the degree of aliasing improves the SNR.²⁹ However, aliasing voxels just beyond the object border can have considerable PSF side lobes that penetrate into the reduced FOV of the object. If a voxel has one of its aliasing peaks just outside the object border, the signal contribution from a PSF side lobe of that peak may be larger than in conventional CSI and may lead to visible artifacts, especially in the presence of strong subcutaneous lipid signal at the edge of the head in brain CSI. Therefore, aliasing voxels lying within one side lobe of the reduced-FOV border should also be reconstructed. This can be achieved by extrapolating the sensitivity maps beyond the object.

The SNR in accelerated parallel imaging is always lower compared to fully sampled acquisitions by the square root of the acceleration factor R because fewer data are acquired

(see *The Basics*). In addition, noise may be further increased because of bad conditioning of the inverse problem because the hybrid encoding functions [the S terms in equation (12)] are not orthogonal.²⁸ The propagation of noise in the reconstruction can be calculated using the sensitivity maps and the noise covariance matrix Ψ for each voxel. This local noise enhancement has been expressed as a geometry factor (g -factor).²⁹ The g -factor increases as the overdetermination of equation (12) decreases, for example, with acceleration factors close to or exceeding the number of coils or when coil sensitivity profiles are too similar. The SNR for SENSE at a voxel ρ relative to that obtained with full encoding is

$$\text{SNR}_\rho^{\text{SENSE}} = \frac{\text{SNR}_\rho^{\text{full}}}{g_\rho \sqrt{R}}, \quad g_\rho \geq 1 \quad (15)$$

The g -factor varies locally and depends on R , the object and the receiver coils used. The g -factor at ρ can be determined from

$$g_\rho = \sqrt{[(S^H \Psi^{-1} S)^{-1}]_{\rho,\rho} (S^H \Psi^{-1} S)_{\rho,\rho}} \quad (16)$$

With Cartesian sampling, SENSE provides an exact solution to the inverse problem with minimal computational effort. Its efficiency reflects the fact that reconstructing a pixel in the final image involves only one pixel in each aliased single-coil image. SENSE cannot only be applied to conventional phase-encoded CSI but also to other faster acquisition methods such as turbo spin echo CSI and proton echo-planar spectroscopic imaging (PEPSI).^{33,34} Yet, it is not easily applicable to CSI with other nuclei such as ^{31}P CSI due to the low SNR of non- ^1H nuclei and the need for coil sensitivity maps, which basically requires either co-registered computed B_1^+ maps or phantom studies of concentrate.

An example comparing single slice ^1H SENSE-CSI with regular CSI from a healthy brain at 3 T is presented in Figure 13. The CSI scan was acquired in 20 min with a 32×32 matrix, a circular shutter, pre-localization with PRESS and OVS (Figure 13b); the SENSE-CSI was acquired with the same PRESS and OVS settings, plus a SENSE factor of 2 in each phase-encoding dimension, which reduced the scan time to 5 min. While the NAA metabolite maps nearly show no difference (Figure 13d), a loss in SNR commensurate with the shorter scan time is evident in the spectra (Figure 13c and e).

GRAPPA

In contrast to SENSE, GRAPPA reconstructs the undersampled data in k -space. Instead of sensitivity maps, the GRAPPA algorithm needs ACS data, which can be acquired as additional phase-encoding steps at the center of k -space (black filled circles in Figure 12F). ACS data can be measured with any contrast, but it is crucial to use the same FOV as in the fully sampled data. For ^1H CSI the ACS data can be obtained from conventional MRI scans, without adding any phase-encoding steps to the MRS sequence, which would increase scan time.³⁵ The GRAPPA reconstruction process is performed in two steps as follows.

First, the ACS data are used as a training set to find weights $W(o, \text{trg}, p, \text{src})$ to allow for a linear combination of the measured, undersampled k -space points, called source points (src),




Figure 13. Example of a single-slice CSI acquisition (32×32 matrix, TR/TE = 1500/144 ms, circular shutter) with and without SENSE from the same transverse MRI slice (shown in a). (b) NAA map calculated from a simple spectral integration of the NAA peak and (c) spectrum acquired with conventional CSI in 20 min. (D) The NAA map and (e) spectrum corresponding to (c) acquired with a SENSE factor of 2 in both phase-encoding directions, which reduced the acquisition time to 5 min. While the metabolite map in (d) does not show the loss of SNR due to reduction in acquisition time, it is evident in (e). Note that these are magnitude spectra that were processed with a digital shift filter to further reduce the water peak. This introduces a sinusoidal weighting of the spectra, attenuating the choline and creatine peaks relative to NAA

to produce the missing k -space points, called target points (trg), as described by

$$s^{\text{trg}}(o, \text{trg}, \text{rep}) = \sum_p \sum_{\text{src}} W(o, \text{trg}, p, \text{src}) s^{\text{src}}(p, \text{src}, \text{rep}) \quad (17)$$

Here o and p count the receive channels for the target and source signals, respectively. The linear combination is performed over all measured k -space points that are neighbors to the missing points within a so-called GRAPPA kernel (white box in Figure 12F) whose size can be chosen. It is important to realize that the weights W are independent of the k -space locations. Therefore, one can slide the kernel through the whole ACS k -space and solve the linear equation system for all those repetitions of the undersampling pattern [index 'rep' in equation (17)]. The weights W can be computed using the pseudo-inverse of s^{src} in equation (17), as equation (17) can be rewritten as a matrix equation. Once the weights W are known, they can be applied in the second step to the measured data to retrieve the missing ones. This corresponds to solving equation (17) for s^{trg} , one missing target point at a time (gray box in Figure 12F).

Although SENSE and GRAPPA perform similarly, there are two key differences. GRAPPA limits the k -space points that are used for reconstruction to the neighboring ones. This makes GRAPPA an intrinsically regularized reconstruction method, whereas SENSE is the exact solution of the inverse problem, and is thus not regularized. The other main difference is that the SENSE algorithm performs the parallel imaging

reconstruction and the coil combination at once, whereas GRAPPA results in uncombined data, which must then be combined using a phased-array algorithm (see section titled 'CSI with Phased-arrays'). Both SENSE and GRAPPA have been extended to non-Cartesian sampling schemes. However, these reconstructions are much more complicated and beyond the scope of this article.

Clinical Example of CSI

An example of a multislice SENSE-CSI scan at 1.5 T in a patient with multiple brain metastases from a melanoma is shown in Figures 14 and 15. The patient had previously received radiation and palliative therapy; however, more metastases had subsequently appeared, which had not been treated with radiation therapy at the time of the MRS. The CSI scan was performed to check the response of the metastases to treatment and was planned to cover three brain slices with a matrix of 32×32 voxels each. OVS alone was used to suppress the subcutaneous lipids in order not to miss metabolite signals in the more peripheral cortical areas. This example illustrates several key aspects of CSI discussed in this article.

First, without some acceleration of the scan time, not all metastases could be observed in one CSI scan because the scan time would have taken 58 min, for which no patient could lie completely still. The SENSE approach, with a SENSE factor of 2 in each phase-encoding dimension, in combination with a




Figure 14. Multislice SENSE-CSI acquisition at 1.5 T of three transverse brain slices in a patient with metastasis from a melanoma. A Cartesian CSI acquisition scheme (32×32 matrix per slice, $TR/TE = 1500/144$ ms, nominal resolution $7 \times 7 \times 15$ mm³) with a SENSE factor of 2×2 , circular shutter, and outer volume suppression slabs was used. The total acquisition time was 14 min. The four columns from left to right show the spin–spin relaxation (T_2)-weighted MRI sections, the NAA, (total) creatine (Cre), and choline (Cho) maps for all three slices, respectively. The resolution of this CSI acquisition suffices to distinguish metastases previously treated with radiation therapy that show necrosis and edema (hyperintense on T_2 -weighted MRI and signal voids on NAA, Cre, and Cho maps), from those showing active cell growth as evidenced by elevated Cho signal (black arrows in slices 2 and 3), sometimes with a necrotic core (slice 3). The hyperintense signal in all three metabolite maps in slice 1 is due to bleeding of the frontal metastasis in slice 1, causing local failure of the water suppression

circular shutter, reduced the scan time for a multislice CSI scan to a tolerable scan time of 14 min.

Second, a nominal spatial resolution of $7 \times 7 \times 15$ mm³ was sufficient to distinguish those metastases that had already received radiation treatment and mainly show necrosis and edema (as evidenced by hyperintensity in the spin–spin relaxation (T_2)-weighted images and a signal void in the NAA, creatine and choline maps) as compared to those showing active cell growth, as indicated by an elevated choline signal (black arrows in slices 2 and 3 of Figure 14), and the suggestion of a necrotic center (slice 3).

Third, using only OVS to suppress subcutaneous lipids, significant residual lipid signal is evident in these integrated metabolite maps. Yet, the lipid signal is suppressed enough to render it comparable to the metabolites of interest, so that it does not interfere appreciably with their interpretation and quantification. In addition, this approach enabled the collection of metabolite signal from peripheral cortical areas.

Fourth, the WET water suppression used for this acquisition performed reasonably well in most regions. However, bleeding in one of the metastases (slice 1), which is typical for melanoma, introduced strong local B_0 inhomogeneity that caused the water suppression to fail locally, giving rise to artifacts in the metabolite maps. Depending on the post-processing methods used to generate the metabolite maps (i.e., integration versus fitting, baseline corrections, etc.), the underlying water peak can show up as a hyperintense signal in all of the metabolite peaks, as it does here. While the choline maps from this patient show well-defined regions of elevated choline, this artifact demonstrates that metabolite maps should be used to guide, rather than replace, a careful inspection of the spectra for clinical purposes.

Fifth, an TE of 144 ms was chosen here to permit easy differentiation of lactate and lipid signals. Figure 15 shows the spectra acquired within the yellow box. The center of the metastasis shows only an inverted lactate peak. The surrounding voxels




Figure 15. Sixteen spectra from the yellow outlined box in the image (a) of slice 2 in Figure 14 (a patient with brain metastases) are displayed (b). In the center of the lesion (bottom right corner) an inverted lactate peak (at $TE = 144$ ms) is evident, along with high choline relative to the missing NAA and creatine, whose intensity increases in the voxels surrounding the lesion, transitioning to a normal appearance in uninvolved brain in the top left corner

show a relatively elevated choline peak relative to the missing NAA and creatine signals, whereas the opposite corner shows nearly normal brain spectra.

Biographical Sketches

Michael Schär obtained a diploma in physics in 2001, and a PhD in biomedical engineering in 2005, both at the Swiss Federal Institute of Technology in Zürich, Switzerland. He was a senior clinical scientist with Philips Healthcare, 2005–2014. Currently, Michael is a research associate at the Johns Hopkins University. He has contributed 66 peer-reviewed publications, 2 patents, and over 120 published abstracts. His research interests are cardiac MRI and MRS, including localized precursors for more robust acquisitions at high field strength.

Bernhard Strasser obtained a master's degree in physics in 2012 at the University of Vienna, Austria. He is currently enrolled as a PhD student at the Vienna University of Technology in physics, and at the Medical University of Vienna in medical physics. He has contributed six peer-reviewed publications. His fields of research are parallel imaging in CSI, fast CSI sequences, and CSI at ultra-high magnetic fields.

Ulrike Dydak obtained her diploma in physics from the University of Vienna, Austria, in 1996, post-graduate diploma in medical physics (2000), and PhD in biomedical engineering (2002) from the Swiss Federal Institute of Technology (ETH) in Zürich, Switzerland. Currently, Ulrike is Associate Professor and Director of the Life Science MRI Facility at Purdue University, with adjunct appointments at the Department of Radiology and Imaging Sciences, Indiana University School of Medicine, and the Department of Biomedical Engineering, Indiana University-Purdue University Indianapolis. She has contributed 38 peer-reviewed publications, 1 patent, 1 book chapter, over 150 published abstracts, and has received the Outstanding New Environmental Scientist award from NIEHS/NIH. Her research interests are MRI and MRS of neurodegenerative and movement disorders and their relation to occupational or environmental exposure to neurotoxins, with special emphasis on GABA-editing and fast spectroscopic imaging.

References

1. J. J. Ackerman, T. H. Grove, G. G. Wong, D. G. Gadian, and G. K. Radda, *Nature*, 1980, **283**, 167.
2. B. Chance, S. Eleff, and J. S. Leigh, *Proc. Natl. Acad. Sci. U. S. A.*, 1980, **77**, 7430.
3. B. Chance, S. Eleff, J. S. Leigh, D. Sokolow, and A. Sapega, *Proc. Natl. Acad. Sci. U. S. A.*, 1981, **78**, 6714.
4. P. A. Bottomley, *Ann. N. Y. Acad. Sci.*, 1987, **508**, 333.
5. J. Frahm, K. D. Merboldt, and W. Hänicke, *J. Magn. Reson.*, 1987, **72**, 502.
6. R. Ordidge, A. Connelly, and J. A. Lohman, *J. Magn. Reson.*, 1986, **66**, 283.
7. W. A. Edelstein, J. M. Hutchison, G. Johnson, and T. Redpath, *Phys. Med. Biol.*, 1980, **25**, 751.
8. D. L. Parker, *Med. Phys.*, 1987, **14**, 640.
9. J. G. Pipe, *Magn. Reson. Med.*, 2000, **43**, 867.
10. A. Greiser and M. von Kienlin, *Magn. Reson. Med.*, 2003, **50**, 1266.
11. J. I. Jackson, C. H. Meyer, D. G. Nishimura, and A. Macovski, *IEEE Trans. Med. Imaging*, 1991, **10**, 473.
12. B. S. Li, J. Regal, and O. Gonen, *Magn. Reson. Med.*, 2001, **46**, 1049.
13. R. Ordidge and I. Cresshull, *J. Magn. Reson.*, 1986, **69**, 151.
14. U. Klose, *Magn. Reson. Med.*, 1990, **14**, 26.
15. P. A. Bottomley, T. H. Foster, and W. M. Leue, *Lancet Lond. Engl.*, 1984, **1**, 1120.
16. A. Haase, J. Frahm, W. Hänicke, and D. Matthaei, *Phys. Med. Biol.*, 1985, **30**, 341.
17. R. J. Ogg, P. B. Kingsley, and J. S. Taylor, *J. Magn. Reson. B*, 1994, **104**, 1.
18. I. Tkáč, Z. Starcuk, I. Y. Choi, and R. Gruetter, *Magn. Reson. Med.*, 1999, **41**, 649.
19. H. Zhu, R. Ouwerkerk, and P. B. Barker, *Magn. Reson. Med.*, 2010, **63**, 1486.
20. P. B. Roemer, W. A. Edelstein, C. E. Hayes, S. P. Souza, and O. M. Mueller, *Magn. Reson. Med.*, 1990, **16**, 192.

21. M. A. Brown, *Magn. Reson. Med.*, 2004, **52**, 1207.
22. Z. Dong and B. Peterson, *Magn. Reson. Imaging*, 2007, **25**, 1148.
23. B. Strasser, M. Chmelik, S. D. Robinson, G. Hangel, S. Gruber, S. Trattnig, and W. Bogner, *NMR Biomed.*, 2013, **26**, 1796.
24. T. Prock, D. J. Collins, A. S. K. Dzik-Jurasz, and M. O. Leach, *Phys. Med. Biol.*, 2002, **47**, N39.
25. T. Schäffter, P. Börnert, C. Leussler, I. C. Carlsen, and D. Leibfritz, *Magn. Reson. Med.*, 1998, **40**, 185.
26. P. B. Barker, B. J. Soher, S. J. Blackband, J. C. Chatham, V. P. Mathews, and R. N. Bryan, *NMR Biomed.*, 1993, **6**, 89.
27. P. Christiansen, O. Henriksen, M. Stubgaard, P. Gideon, and H. B. W. Larsson, *Magn. Reson. Imaging*, 1993, **11**, 107.
28. K. P. Pruessmann, *NMR Biomed.*, 2006, **19**, 288.
29. K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger, *Magn. Reson. Med.*, 1999, **42**, 952.
30. M. A. Griswold, P. M. Jakob, R. M. Heidemann, M. Nittka, V. Jellus, J. Wang, B. Kiefer, and A. Haase, *Magn. Reson. Med.*, 2002, **47**, 1202.
31. U. Dydak, M. Weiger, K. P. Pruessmann, D. Meier, and P. Boesiger, *Magn. Reson. Med.*, 2001, **46**, 713.
32. M. Weiger, K. P. Pruessmann, and P. Boesiger, *MAGMA*, 2002, **14**, 10.
33. U. Dydak, K. P. Pruessmann, M. Weiger, J. Tsao, D. Meier, and P. Boesiger, *Magn. Reson. Med.*, 2003, **50**, 196.
34. F.-H. Lin, S.-Y. Tsai, R. Otazo, A. Caprihan, L. L. Wald, J. W. Belliveau, and S. Posse, *Magn. Reson. Med.*, 2007, **57**, 249.
35. F. Breuer, D. Ebel, J. Ruff, M. Blaimer, N. Seiberlich, M. A. Griswold, P. M. Jakob, In *Proceedings of the ISMRM Annual Meeting, Seattle, Washington, USA*, 2006, 3653.