

Perspectives in Magnetic Resonance

Magnetic resonance spectroscopic imaging at superresolution: Overview and perspectives



Jeffrey Kasten^{a,b,*}, Antoine Klauer^b, François Lazeyras^b, Dimitri Van De Ville^{a,b}

^a Institute of Bioengineering, École Polytechnique Fédérale de Lausanne, Switzerland

^b Department of Radiology and Medical Informatics, University of Geneva, Switzerland

ARTICLE INFO

Article history:

Received 9 September 2015

Revised 7 November 2015

Available online 2 December 2015

Keywords:

Magnetic resonance spectroscopic imaging

Superresolution

Chemical shift imaging

Parametric modeling

Regularization

Non-Fourier reconstruction

Denoising

Spatio-spectral deconvolution

Metabolite mapping

ABSTRACT

The notion of non-invasive, high-resolution spatial mapping of metabolite concentrations has long enticed the medical community. While magnetic resonance spectroscopic imaging (MRSI) is capable of achieving the requisite spatio-spectral localization, it has traditionally been encumbered by significant resolution constraints that have thus far undermined its clinical utility. To surpass these obstacles, research efforts have primarily focused on hardware enhancements or the development of accelerated acquisition strategies to improve the experimental sensitivity per unit time. Concomitantly, a number of innovative reconstruction techniques have emerged as alternatives to the standard inverse discrete Fourier transform (DFT). While perhaps lesser known, these latter methods strive to effect commensurate resolution gains by exploiting known properties of the underlying MRSI signal in concert with advanced image and signal processing techniques. This review article aims to aggregate and provide an overview of the past few decades of so-called “superresolution” MRSI reconstruction methodologies, and to introduce readers to current state-of-the-art approaches. A number of perspectives are then offered as to the future of high-resolution MRSI, with a particular focus on translation into clinical settings.

© 2015 Elsevier Inc. All rights reserved.

1. Introduction

Magnetic resonance spectroscopic imaging (MRSI) is a unique molecular imaging modality that offers simultaneous spatial and spectral localization, and has become a powerful clinical research tool. The potential functional or pathological insights that may be obtained, perhaps even prior to the presentation of any structural or behavioral aberrations, through non-invasive spatial mapping of metabolite concentrations has long captivated the clinical research community. Yet despite its allure, MRSI has been traditionally encumbered by a number of practical limitations. The foremost impediment stems from the relatively low tissue metabolite concentrations, producing signals that are often vastly overwhelmed by those originating from the primary solvent or chemical background. Furthermore, the need to encode the MR signal along an additional temporal dimension precludes the use of many efficient acquisition schemes normally exploited in structural MR imaging, and leads to greatly protracted exam durations. Taken in conjunction, these considerations ultimately necessitate compromises in the data collection process that may carry significant consequences with respect to the final reconstructed data.

In most clinical settings, MRSI data are acquired on a rectilinear k -space grid, where typically only low spatial frequency information (which generally accounts for the greatest portion of the total signal energy) is encoded in order to maximize the experimental sensitivity while minimizing the overall exam time. This basic acquisition approach is also commonly referred to as a chemical shift imaging (CSI) experiment, and will be referred to as such for the remainder of the text. The data are then customarily reconstructed by means of the inverse discrete Fourier transform (DFT). While this type of standardized acquisition/reconstruction protocol is often pragmatic, the resultant voxel sizes following DFT are typically on the order of 0.5–2.0 cm³. Such coarse resolutions preclude the ability to identify subtle spatial features in the spectroscopic signal, thereby limiting the overall impact of MRSI. Moreover, this strategy necessarily entails a broad system point-spread function (PSF), precipitating so-called spectral leakage effects whereby local spectra are contaminated by signal contributions originating from spatially remote regions, for example, the extra-cranial lipids in ¹H studies of the brain. Although standard pre-processing pipelines often include additional k -space filtering steps to mitigate these effects, such methods tend to incur further losses in an already insufficient spatial resolution. Another common approach is to refrain from exciting/refocusing the signal from known “nuisance” components during the acquisition through

* Corresponding author at: Department of Radiology and Medical Informatics, University of Geneva, Switzerland.

localization methods such as STEAM [1–3] or PRESS [4–7]. However, such techniques preclude the use of non-rectangular volumes of interest (VOIs), which may consequently exclude portions of the anatomy of interest as well.

Over the past few decades, concerted efforts have been made in order to surmount the aforementioned limitations that have thus far encumbered MRSI. Although by no means mutually exclusive, these endeavors can typically be characterized as falling into one of two categories. The first pertains to the acquisition process, where a number of accelerated schemes have been developed in order to probe more distal regions in k -space while maintaining comparable signal-to-noise ratios (SNR) and sensitivities per unit time to traditional encoding strategies. The second revolves around the reconstruction procedure itself, whereby alternative methods have been sought that aim to circumvent the resolution and geometry restrictions imposed through direct use of the DFT. While the former have been mostly circumscribed by MR-specific literature (for reviews, see [8,9]), descriptions of the latter have been somewhat fragmented throughout the biomedical imaging, signal processing, as well as MR communities. The aim of this review is to therefore familiarize readers with these so-called “superresolution” reconstruction methods, providing some historical background while focusing on current state-of-the-art approaches. We then conclude with some perspectives on the various difficulties that have thus far impeded the translation of such endeavors into routine clinical settings.

2. What is meant by superresolution reconstruction?

Superresolution in MRSI can essentially be described as the concept of reconstructing a dataset at a higher spatial resolution than that which would be typically dictated by the nominal measurement voxel size, while simultaneously achieving a similar or higher SNR. However, contrary to conventional notions of superresolution, whereby ancillary encoding techniques are employed in order to transcend fundamental limits imposed by the system PSF, such enhancements are achieved within the context of the current survey primarily through the use of sophisticated computational methods that exploit expected signal characteristics and known system properties. This could cover, for instance, tailored denoising schemes that attempt to improve the SNR of high-resolution MRSI datasets achieved through accelerated acquisitions, or simply attempts to generate high-resolution spatial metabolite profiles in the wake of limited k -space data, thereby defining a classical “inverse problem”. Regardless, such issues are generally approached by introducing some type of *a priori* knowledge pertaining to the acquisition process and/or the imaged object itself, which forms the basis for a corresponding signal model. Such prior information may stem from properties ascertained through attendant acquisitions, the set of physical principles governing the measurement process, or may simply reflect a wealth of qualitative observations and precepts obtained through years of clinical experience. Whatever the case may be, the fundamental challenge for any prospective reconstruction scheme is to successfully translate such knowledge into a quantitative and robust framework. The following sections describe the primary foundations and principles underlying superresolution MRSI reconstruction, highlighting some of the key developments leading up to current state-of-the-art methods. For preliminary definitions and notations, readers are referred to Table 1.

3. The model-based reconstruction framework

Disregarding relaxation effects, as well as added experimental confounds such as noise or spatially-dependent inhomogeneities

Table 1

Various mathematical notations used throughout the manuscript.

j	Imaginary unit such that $j^2 = -1$
z^*	Complex conjugate of the complex number z
$\mathbb{1}_C$	Characteristic function of the set C
\mathcal{P}_C	Euclidean projector onto the set C
\mathbf{X}^H	Hermitian transpose of the matrix \mathbf{X}
$ \cdot $	Absolute value of the argument
$\ f\ _{L^p}$	L^p -norm of a continuous-domain function, $\ f\ _{L^p} = \begin{cases} (\int_{\Omega} f ^p dm)^{1/p} & 1 \leq p < \infty \\ \sup_{\mathbf{x} \in \Omega} f(\mathbf{x}) & p = \infty \end{cases}$
$\ \mathbf{x}\ _{\ell_p}$	ℓ_p -norm of a vector, $\ \mathbf{x}\ _{\ell_p} = \begin{cases} (\sum_{i=1}^N x_i ^p)^{1/p} & 1 \leq p < \infty \\ \max(x_1 , \dots, x_N) & p = \infty \end{cases}$
$\ \mathbf{X}\ _F$	Frobenius norm of a matrix, $\ \mathbf{X}\ _F = (\sum_{ij} X_{ij} ^2)^{1/2}$
$\ \mathbf{X}\ _S$	Absolute sum norm of a matrix, $\ \mathbf{X}\ _S = (\sum_{ij} X_{ij})$

in the local magnetic field or radiofrequency profile, the MRSI signal can be expressed as:

$$s(\mathbf{k}, t) = \int_{-\infty}^{\infty} \int_{\Omega \subset \mathbb{R}^3} \rho(\mathbf{r}, f) e^{-2\pi j(\mathbf{k} \cdot \mathbf{r} + ft)} d\mathbf{r} df, \quad (1)$$

where ρ is the spatio-spectral distribution of the object, and Ω its spatial support. As with any digital imaging system, however, the MR scanner does not furnish a continuous stream of observations, but rather samples the observed signal and stores a collection of discrete measurements. The practicable discrete data can therefore be expressed as:

$$\tilde{\mathbf{S}} = \mathcal{A}\{\rho\} + \mathbf{N} \quad (2)$$

$$= \mathbf{S} + \mathbf{N}, \quad (3)$$

where \mathcal{A} is a continuous-to-discrete linear mapping that represents the acquisition scheme – consisting in this case of a forward Fourier transform evaluated over the set $\mathcal{X} = \mathcal{K} \otimes \mathcal{T}$, where $\mathcal{K} = \{\mathbf{k}_m\}_{m=1}^M$ and $\mathcal{T} = \{t_i\}_{i=1}^T$ denote the prescribed k -space and temporal sampling locations, respectively – and \mathbf{N} represents additive noise originating primarily from the coil and scanner electronics. Here, $\tilde{\mathbf{S}}, \mathbf{S}, \mathbf{N} \in \mathbb{C}^{M \times T}$ are given as complex-valued $M \times T$ matrices, where each row represents the observed temporal signal corresponding to a particular k -space location, \mathbf{k}_m .

On a fundamental level, the aim of the reconstruction is to obtain a faithful estimate of ρ from the observed (noisy) measurements. As the noise distribution of standard quadrature-detected MR signals is generally assumed to be zero-mean additive white Gaussian (AWGN), i.e., $p(\mathbf{N}) \sim \mathcal{N}(0, \sigma^2)$, where σ^2 is the noise variance, estimates of ρ in (2) could be found through a maximum likelihood framework as:

$$\begin{aligned} \hat{\rho} &= \arg \max_{\rho} p(\mathcal{A}\{\rho\} + \mathbf{N} | \tilde{\mathbf{S}}) \quad (4) \\ &= \arg \max_{\rho} p(\mathbf{N}) \\ &= \arg \max_{\rho} \left(\frac{1}{2\pi\sigma^2} \right)^{1/2} e^{-\frac{\|\tilde{\mathbf{S}} - \mathcal{A}\{\rho\}\|_F^2}{2\sigma^2}} \\ &= \arg \max_{\rho} \log \left(\left(\frac{1}{2\pi\sigma^2} \right)^{1/2} e^{-\frac{\|\tilde{\mathbf{S}} - \mathcal{A}\{\rho\}\|_F^2}{2\sigma^2}} \right) \\ &= \arg \max_{\rho} \left(-\frac{\|\tilde{\mathbf{S}} - \mathcal{A}\{\rho\}\|_F^2}{2\sigma^2} \right) + C \\ &= \arg \min_{\rho} \|\tilde{\mathbf{S}} - \mathcal{A}\{\rho\}\|_F^2. \quad (5) \end{aligned}$$

In essence, the above procedure amounts to finding an estimate of ρ that, under the model and noise assumptions, yields a set of

theoretical measurements that are maximally consistent with those acquired in practice. Expression (5) itself is often referred to as the “least squares” (LS) or “linear regression” problem, and is often regarded as the “variational formulation” within the wider study of inverse problems.

Unfortunately, the integral defined by $\mathcal{A}\{\rho\}$ in (2) can seldom be evaluated analytically for arbitrary object geometries, and must therefore be approximated numerically. This is commonly accomplished by considering the fully-discretized expression:

$$\tilde{\mathbf{s}}[t_i] = \mathbf{A}\mathbf{x}[t_i] + \mathbf{n}[t_i]. \quad (6)$$

Here, $\mathbf{A} \in \mathbb{C}^{M \times N}$ represents the discretized encoding scheme:

$$\mathbf{A} = \begin{pmatrix} e^{-2\pi j \mathbf{k}_1 \cdot \mathbf{r}_1} & e^{-2\pi j \mathbf{k}_1 \cdot \mathbf{r}_2} & \dots & e^{-2\pi j \mathbf{k}_1 \cdot \mathbf{r}_N} \\ e^{-2\pi j \mathbf{k}_2 \cdot \mathbf{r}_1} & e^{-2\pi j \mathbf{k}_2 \cdot \mathbf{r}_2} & \dots & e^{-2\pi j \mathbf{k}_2 \cdot \mathbf{r}_N} \\ \vdots & \vdots & \ddots & \vdots \\ e^{-2\pi j \mathbf{k}_M \cdot \mathbf{r}_1} & e^{-2\pi j \mathbf{k}_M \cdot \mathbf{r}_2} & \dots & e^{-2\pi j \mathbf{k}_M \cdot \mathbf{r}_N} \end{pmatrix}, \quad (7)$$

where $\{\mathbf{r}_n\}_{n=1}^N$ denotes the set of sampling points in the spatial domain, while,

$$\mathbf{x}[t_i] = \begin{pmatrix} x_1[t_i] \\ x_2[t_i] \\ \vdots \\ x_N[t_i] \end{pmatrix}, \quad (8)$$

contains the corresponding signal amplitude coefficients at each temporal sample point, t_i . For completeness, $\tilde{\mathbf{s}}[t_i]$ and $\mathbf{n}[t_i]$ represent individual columns of $\tilde{\mathbf{S}}$ and \mathbf{N} in (2), respectively. The resulting reconstruction problem is then given by:

$$\hat{\mathbf{x}}[t_i] = \arg \min_{\mathbf{x}[t_i]} \|\tilde{\mathbf{s}}[t_i] - \mathbf{A}\mathbf{x}[t_i]\|_{\ell_2}^2, \quad (9)$$

which admits a closed-form solution:

$$\hat{\mathbf{x}}[t_i] = (\mathbf{A}^H \mathbf{A})^{-1} \mathbf{A}^H \tilde{\mathbf{s}}[t_i] = \mathbf{A}^+ \tilde{\mathbf{s}}[t_i], \quad (10)$$

where $(\cdot)^+$ denotes the Moore–Penrose pseudoinverse.

4. Model-based reinterpretation of the DFT

In the case where (7) is square ($M = N$), the matrix inversion in (10) reduces to identity, thereby reproducing the standard inverse DFT reconstruction. While a familiar and prominent tool, additional insights surrounding the DFT may be gained through the above model-based formulation that carry a number of implications for MRSI reconstruction. First, although the DFT is often regarded as an unbiased and generalized reconstruction procedure, it does prescribe an implicit object model. This can be seen more clearly by rewriting (6) as:

$$\tilde{\mathbf{s}}[t_i] = \mathbf{A}(\mathbf{A}^H \mathbf{s}[t_i]) + \mathbf{n}[t_i]. \quad (11)$$

In other words, the DFT implicitly posits that the underlying object is strictly band-limited (and periodic). A second and closely related observation is that the precise “band-limitedness” of the model, and consequently, the resulting reconstruction grid, is ultimately dictated by \mathcal{K} . Therefore, as the measurement protocol becomes more stringent and fewer k -space encodings are permitted, the DFT model will expectedly represent the underlying object by increasingly coarse approximations (i.e., at lower spatial resolution). Clearly, a much more accurate description of the measurement process would follow by effectively decoupling the acquisition and reconstruction grids, allowing $N \gg M$ in (7), which can be considered as the crux of “superresolution” reconstruction. This actualized as such, however, would customarily endow \mathbf{A} with a large null

space, making the reconstruction in (10) ill-posed. To counter this limitation, preliminary attempts at “superresolution” proceed by appending the data with zeros prior to DFT, thereby artificially inflating M up to the desired N . Commonly known as the “zero-filled DFT solution” (ZDFT), this procedure generally produces more visually appealing results (see Fig. 1), but does not introduce any additional information content, and merely reflects spatial interpolation of the low resolution data by the system PSF (the Dirichlet kernel), which is fundamentally determined by \mathcal{K} .

5. Reconstruction by Explicit Parametric Modeling

In light of the above discussion, it is natural to question whether alternative parametrizations of the underlying signal, especially those that exploit prior object knowledge, may be elected over the discrete Fourier representation intimated by (11). Indeed, although MRSI represents an ultimately unique modality, certain properties of the MRSI signal may be shared among other modalities for which resolution and SNR constraints are not as stringent. One of the earliest methods to explore these relationships was the spectral localization by imaging (SLIM) technique [10]. Standing in the vanguard of so-called “constrained reconstruction” approaches, SLIM postulated that the spatial distribution of chemical species of interest tends to parallel that of water in biological samples. Therefore, knowledge of anatomical features derived from high resolution structural ^1H MRI images – which are typically acquired as part of any standard scanning regime – could be used as *a priori* information during spectroscopic image reconstruction. In a typical SLIM experiment, the underlying object is modeled as a superposition of K anatomically-defined compartments, which are assumed to be spectrally-homogeneous, such that:

$$\rho_{slim}(\mathbf{r}, f) = \sum_{\kappa=1}^K \mathbb{1}_{\kappa}(\mathbf{r}) q_{\kappa}(f), \quad (12)$$

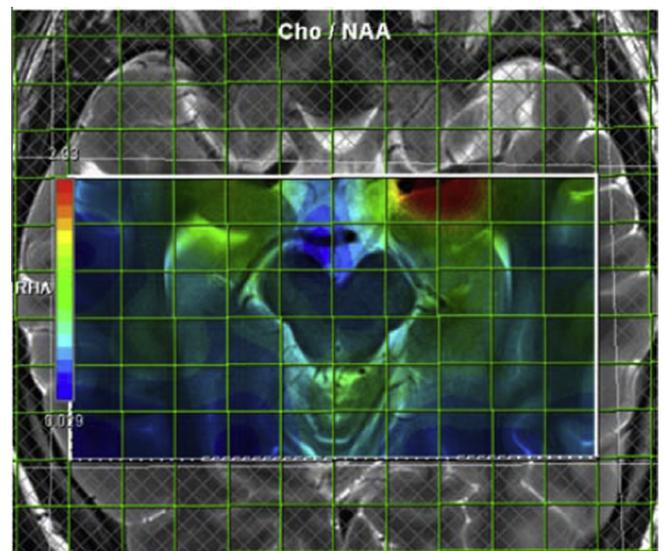


Fig. 1. Typical MRSI data acquisition and display. Here, the MRSI data (delimited by the outer white box marking the excitation volume) has been overlaid on a high-resolution structural image acquired during the same scanning session. The false color image represents a ZDFT reconstruction visualizing the spatial distribution of the choline (Cho)/N-Acetylaspartic acid (NAA) ratio. The nominal in-plane voxel dimensions, however, are delineated by the solid green lines.

where q_κ is the unknown spectrum, and $\mathbb{1}_\kappa$ the characteristic function for the κ th compartment:

$$\mathbb{1}_\kappa(\mathbf{r}) \triangleq \begin{cases} 1 & \mathbf{r} \in \text{compartment } \kappa \\ 0 & \text{otherwise.} \end{cases} \quad (13)$$

Disregarding noise, and substituting (12) for ρ in (1), the unperturbed (i.e., theoretically predicted) MRSI signal can be expressed as:

$$\begin{aligned} s(\mathbf{k}, t) &= \int_{-\infty}^{\infty} \int_{\Omega \subset \mathbb{R}^3} \sum_{\kappa=1}^K \mathbb{1}_\kappa(\mathbf{r}) q_\kappa(f) e^{-2\pi j(\mathbf{k}\mathbf{r}+ft)} d\mathbf{r} df \\ &= \sum_{\kappa=1}^K \underbrace{\int_{\Omega \subset \mathbb{R}^3} \mathbb{1}_\kappa(\mathbf{r}) e^{-2\pi j\mathbf{k}\mathbf{r}} d\mathbf{r}}_{X_\kappa(\mathbf{k})} \cdot \underbrace{\int_{-\infty}^{\infty} q_\kappa(f) e^{-2\pi jft} df}_{c_\kappa(t)} \\ &= \sum_{\kappa=1}^K X_\kappa(\mathbf{k}) c_\kappa(t). \end{aligned} \quad (14)$$

Echoing the previous discussion, while the theoretical signal is here presented in the continuous domain (as will generally be the case for the remainder of the text), the integral expressions in (14) can be seldom evaluated analytically, leading to the following discretization:

$$\mathbf{s}[t_i] = \mathbf{X}\mathbf{c}[t_i], \quad (15)$$

where

$$\mathbf{X} = \begin{pmatrix} X_1[\mathbf{k}_1] & X_2[\mathbf{k}_1] & \cdots & X_K[\mathbf{k}_1] \\ X_1[\mathbf{k}_2] & X_2[\mathbf{k}_2] & \cdots & X_K[\mathbf{k}_2] \\ \vdots & \vdots & \ddots & \vdots \\ X_1[\mathbf{k}_M] & X_2[\mathbf{k}_M] & \cdots & X_K[\mathbf{k}_M] \end{pmatrix}. \quad (16)$$

In other words, the specified object model effects a change of basis from \mathbf{A} (see (7)) to \mathbf{X} for the observed signal. The corresponding amplitude coefficients at each time point, $\mathbf{c}[t_i]$, can then be estimated using the standard LS scheme as:

$$\hat{\mathbf{c}}[t_i] = \arg \min_{\mathbf{c}[t_i]} \|\tilde{\mathbf{s}}[t_i] - \mathbf{X}\mathbf{c}[t_i]\|_{\ell_2}^2. \quad (17)$$

When contrasting (16) and its DFT counterpart ((7) with $M = N$), a number of important distinctions can be made. First and foremost is that unlike the DFT, the number of basis elements (columns of (16)) remain independent of the number of encoding steps. In other words, reconstructed volumes are no longer circumscribed by the nominal voxel dimensions as dictated by the collected k -space measurements. Rather, the compartments can be said to define so-called “generalized voxels”, the geometries of which can be adapted to local anatomical features and discretized on the same high-resolution grids as those realized by structural imaging. Secondly, the linear system defined by (15) remains over-determined as long as the number of compartments does not exceed the number of k -space measurements, thereby raising the prospect of estimating the associated compartmental spectra from as few as K phase encoding steps. Such reduced sampling requirements could therefore translate to drastically reduced exam times for certain applications where the compartmental model is deemed appropriate. Other characteristics of the SLIM framework, which are discussed in depth in [11,12], are that the noise sensitivity is inversely proportional to compartmental volume, and that estimated spectra converge to the true compartmental average as k -space coverages increases.

Since SLIM's inception, a number of extensions have been proposed that exploit the various advantages afforded by the basic framework, some of which have been previously discussed in an early review by Liang et al. [13], but which are included here for convenience. For example, although the theory states that in

principle only K encoding steps are necessary to ensure a unique solution to (17), certain choices for K may be more judicious than others. The spectral localization with optimal pointspread function (SLOOP) method [14] suggested a means for optimizing experimental sensitivity while minimizing the potential for signal contamination in a SLIM experiment. This was accomplished by establishing a criterion for selecting K so as to tailor the spatial response function (SRF) to the geometry of each compartment, where:

$$\text{SRF}_\kappa(\mathbf{r}) = \sum_{m=1}^M X_\kappa^+[\mathbf{k}_m] e^{-2\pi j\mathbf{k}_m \cdot \mathbf{r}}. \quad (18)$$

Here, the SRF represents the spatial origin of all contributions to the estimated signal associated with a particular basis element (i.e., the compartments in the SLIM framework). Similar efforts include that of [15], which sought to optimize K by minimizing the expected LS error given the image support region. Another key advantage of the SLIM framework's dissociation from the elected k -space samples is that it remains fully compatible with the bevy of accelerated acquisition techniques that have been developed contemporaneously. For example, [16] assessed the spectral contents of small lesions in stroke patients by combining SLIM with sensitivity encoding (SENSE) [17], allowing for even further reductions in scan time.

One of the most contentious aspects of the SLIM technique, however, is its reliance upon the spectral homogeneity assumption within each compartment. While appropriate or beneficial for certain applications, such expectations are rarely met in common clinical settings. In reality, additional experimental confounds such as field inhomogeneities, inappropriate designations for the number or geometry of the compartments, registration errors, or any other factors that may lead to discrepancies between the signal model and the acquired measurements will engender additional artifacts. In order to add flexibility to the basic model, the GSLIM method [18] was proposed as a compromise between SLIM and traditional Fourier reconstructions, profiting from both the implicit high-resolution information furnished by the former while capturing unanticipated spatial variations through the use of Fourier-type spatial harmonics. In GSLIM, the spatio-spectral distribution is modeled as:

$$\rho_{\text{gslim}}(\mathbf{r}, f) = \sum_{l=1}^L \underbrace{\sum_{\kappa=1}^K \mathbb{1}_\kappa(\mathbf{r}) q_\kappa(f) a_l(f)}_{\rho_{\text{slim}}(\mathbf{r}, f)} e^{2\pi j\mathbf{k}_l \cdot \mathbf{r}}, \quad (19)$$

such that the resulting signal becomes:

$$\begin{aligned} s(\mathbf{k}, t) &= \sum_{l=1}^L \sum_{\kappa=1}^K \int_{\Omega \subset \mathbb{R}^3} \mathbb{1}_\kappa(\mathbf{r}) e^{-2\pi j(\mathbf{k}-\mathbf{k}_l) \cdot \mathbf{r}} d\mathbf{r} \\ &\quad \cdot \int_{-\infty}^{\infty} a_l(f) q_\kappa(f) e^{-2\pi jft} df. \end{aligned} \quad (20)$$

As q_κ and $\mathbb{1}_\kappa$ are pre-determined using the original SLIM framework, the GSLIM coefficients, a_l – designed to absorb any residual spatially-dependent spectral variations – can be calculated by expressing (20) in the temporal Fourier domain:

$$\mathcal{F}_t\{s\}(\mathbf{k}, f) = \sum_{l=1}^L a_l(f) \underbrace{\sum_{\kappa=1}^K q_\kappa(f) \int_{\Omega \subset \mathbb{R}^3} \mathbb{1}_\kappa(\mathbf{r}) e^{-2\pi j(\mathbf{k}-\mathbf{k}_l) \cdot \mathbf{r}} d\mathbf{r}}_{G(\mathbf{k}-\mathbf{k}_l, f)}, \quad (21)$$

and solving the ensuing (discrete) linear system.

To specifically address the contribution of local (static) field inhomogeneities, the BSLIM method [19] relies on an additional measured field map, $\Delta f(\mathbf{r})$, in order to correct for any concomitant spectral shifts, which manifest as spurious spectral peaks or

distorted lineshapes in the estimated SLIM spectra if unaccounted for. The object model then becomes:

$$\rho_{bslim}(\mathbf{r}, f) = \sum_{\kappa=1}^K \mathbb{1}_{\kappa}(\mathbf{r}) q_{\kappa}(f - \Delta f(\mathbf{r})), \quad (22)$$

leading to a measured signal:

$$\begin{aligned} s(\mathbf{k}, t) &= \int_{-\infty}^{\infty} \int_{\Omega \subset \mathbb{R}^3} \sum_{\kappa=1}^K \mathbb{1}_{\kappa}(\mathbf{r}) q_{\kappa}(f - \Delta f(\mathbf{r})) e^{-2\pi j(\mathbf{k}\mathbf{r} + ft)} d\mathbf{r} df \\ &= \sum_{\kappa=1}^K \underbrace{\int_{\Omega \subset \mathbb{R}^3} \mathbb{1}_{\kappa}(\mathbf{r}) e^{-2\pi j(\mathbf{k}\mathbf{r} + \Delta f(\mathbf{r})t)} d\mathbf{r}}_{H_{\kappa}(\mathbf{k}, t)} \cdot \underbrace{\int_{-\infty}^{\infty} q_{\kappa}(f) e^{-2\pi jft} df}_{c_{\kappa}(t)} \\ &= \sum_{\kappa=1}^K H_{\kappa}(\mathbf{k}, t) c_{\kappa}(t). \end{aligned} \quad (23)$$

It is clear when contrasting (23) with (14) that the principal difference between the methods lies in the acquired time dependence of the basis, H_{κ} , which subsumes the position-dependent phase shifts induced by $\Delta f(\mathbf{r})$. In other words, BSLIM fully subscribes to the compartmental homogeneity precept under the assumption that any intra-compartmental spectral heterogeneities are solely attributable to local field variations. For a visual guide contrasting the various SLIM-type basis functions, see Fig. 2.

Throughout the development of MRSI reconstruction methodologies, SLIM has remained a perennial force owing to its conceptual ease and extensibility. In addition to the above examples, SLIM and/or its variants have also been utilized for ^1H applications using pre-clinical models of stroke [20], post-acquisition suppression of extramyocellular lipids in the human calf [21], as well as for ^{31}P studies of the human myocardium [22–24]. Nevertheless, SLIM-type methods remain reliant upon explicit knowledge of the underlying compartments, and are therefore fundamentally circumscribed by the accuracy of the spatial priors. Furthermore,

general difficulties in obtaining accurate segmentation data, let alone possible inaccuracies in the basic presumption that metabolite distributions parallel that of the structural anatomy, have thus far limited SLIM’s utility in general clinical MRSI applications.

6. Reconstruction via regularized approaches

The aforementioned concerns surrounding reconstruction bias in the face of inaccurate or overly-constraining priors prompted the reconstruction community to consider whether such knowledge could be incorporated into the reconstruction procedure in a “softer” or more flexible manner. Out of this dialogue grew an increasing number of regularized methods, which allowed greater control over the influence of the selected priors. Mathematically, regularization typically assumes the form of an additional penalty in conjunction with the standard quadratic term in the LS framework:

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \underbrace{\|\tilde{\mathbf{s}} - \mathbf{A}\mathbf{x}\|_{\ell_2}^2}_{\text{data consistency}} + \underbrace{\mathcal{R}(\mathbf{x})}_{\text{regularization}}, \quad (24)$$

where \mathcal{R} (known as the “regularizer”) is a suitable metric for evaluating the chosen constraints. In this case, superresolution is achieved not by explicitly parametrizing the basis functions comprising the object model as detailed in the previous section, but by selecting the regularizer so as to ensure a well-posed reconstruction problem.

One early example using regularization is that of [25], in which the sole *a priori* information is the spatial support region of the entire imaged object. Adapting the well-established Papoulis–Gerchberg (PG) algorithm [26,27], the reconstruction, dubbed as the “finite support solution” amounts to minimizing the total energy added by any extrapolated k -space coefficients, given the high frequency content implicitly carried by the support region:

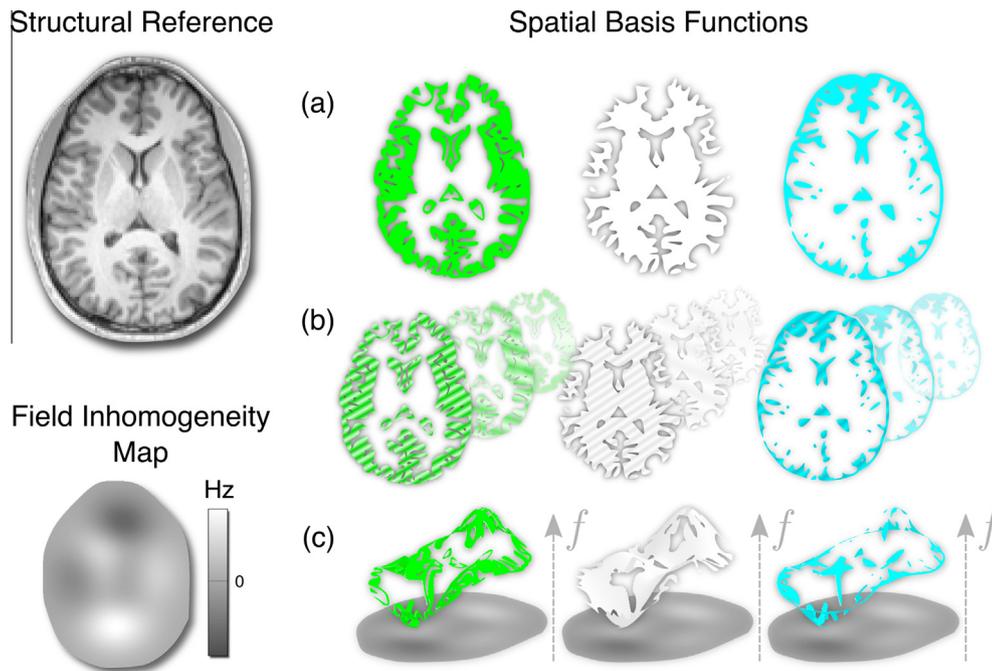


Fig. 2. Example spatial basis functions employed by each of the SLIM variants. (a) Archetypal SLIM basis functions, here consisting of segmented gray matter (GM), white matter (WM), and cerebrospinal fluid (CSF) as derived from accompanying high-resolution structural reference scans (top, left) reflecting the water distribution. (b) GSLIM introduces additional Fourier-type harmonic basis functions, which are confined to the spatial support of the original SLIM basis functions. (c) BSLIM warps the SLIM bases in the temporal frequency domain in accordance with a separately acquired high-resolution local field inhomogeneity map (bottom, left), fitting the unknown spectral components over the resulting manifold.

$$\hat{\mathbf{c}}[f_l] = \arg \min_{\mathbf{c}[f_l]} \|\tilde{\mathbf{s}}[f_l] - \mathbf{F}\mathbf{c}[f_l]\|_{\ell_2}^2 + \lambda \left(\mathbf{c}[f_l]^H \boldsymbol{\Omega} \mathbf{c}[f_l] \right), \quad (25)$$

where $\{f_l\}_{l=1}^L$ is the set of discretized measurements in the temporal frequency domain, $\boldsymbol{\Omega}$ is a binary diagonal matrix specifying the support region, and $\lambda \in \mathbb{R}$ is a parameter mediating the tradeoff between data consistency and regularization bias. In this case, improvements in the resulting PSF were reported throughout the entire reconstructed volume when compared to ZDFT solutions, but were primarily localized to the periphery of the support region (where the prior information is strongest). As with the SLIM and SLOOP methods, the advantages of the finite support solution opened additional avenues for tailored acquisition strategies, for example in [28], where k -space sampling locations were selected so as to minimize leakage artifacts arising from intensity differences between tissue types.

As a natural transition from the SLIM-type frameworks, the method of [29] aimed to exploit the versatility afforded by regularization-based reconstruction while retaining salient object features identified by high-resolution structural MRI scans. In this case, the primary GSLIM model is supplemented by local B-spline basis functions such that:

$$\rho_{imbr}(\mathbf{r}, f) = \underbrace{\sum_{l=1}^L \sum_{\kappa=1}^K \mathbb{1}_{\kappa}(\mathbf{r}) c_{\kappa,l}(f) e^{2\pi j \mathbf{k}_l \cdot \mathbf{r}}}_{\text{compartmental basis functions}} + \underbrace{\sum_{p=1}^{p^2} d_p(f) \beta^3(\mathbf{r} - \mathbf{r}_p)}_{\text{local basis functions}}, \quad (26)$$

where β^3 is a cubic B-spline function. Due to their compact support, such functions maintain a greater capacity for representing local intensity variations that would otherwise require significant Fourier expansions, such as those attributable to small punctate lesions. The signal model, which also accounts for static field inhomogeneities as well as positional inconsistencies between the structural and MRSI data, is then given by:

$$s(\mathbf{k}, t) = \int_{-\infty}^{\infty} \int_{\Omega \subset \mathbb{R}^3} \rho_{imbr}(\mathcal{T}\{\mathbf{r}\} + \mathbf{b}, f) e^{-2\pi j(\mathbf{k} \cdot \mathbf{r} + (f + \Delta f(\mathbf{r}))t)} d\mathbf{r} df, \quad (27)$$

where $\mathcal{T}\{\mathbf{r}\} + \mathbf{b}$ is an affine transformation. The reconstruction is then formulated as the following optimization problem:

$$\hat{\mathbf{c}}[f_l], \hat{\mathbf{d}}[f_l], \hat{\mathbf{T}}, \hat{\mathbf{b}} = \arg \min_{\mathbf{c}[f_l], \mathbf{d}[f_l], \mathbf{T}, \mathbf{b}} \|\tilde{\mathbf{s}}[t_i] - \mathbf{s}[t_i]\|_{\ell_2}^2 + \lambda \left(\|\mathbf{c}[f_l]\|_{\ell_1} + \|\mathbf{d}[f_l]\|_{\ell_1} \right), \quad (28)$$

where \mathbf{T} is the matrix representation of the linear operator comprising the affine transformation. Here, the aim of the regularizer is to promote solutions in which only a small subset of basis functions are “active” at a given temporal frequency location, i.e., $\mathbf{c}[f_l]$ and $\mathbf{d}[f_l]$ are “sparse” vectors in which only a small number of elements are non-zero. Although the most conceptually straightforward means of enforcing sparsity would be through the ℓ_0 pseudo-norm, which returns the number of non-zero coefficients of an input vector, this penalty is fundamentally non-convex. Indeed, problems of the form:

$$\arg \min_{\mathbf{x}} \|\tilde{\mathbf{s}} - \mathbf{A}\mathbf{x}\|_{\ell_2}^2 + \lambda \|\mathbf{x}\|_{\ell_0}, \quad (29)$$

are NP-hard [30], and global minima can only be obtained through computationally demanding combinatorial search methods or heuristic approaches [31,32]. To therefore leverage the wealth of efficient algorithms that have been developed for convex optimization, the ℓ_0 penalty is often supplanted by the ℓ_1 norm as the closest convex approximation. Solutions to (28) can then be found via an

alternating scheme, first estimating the GSLIM and B-spline coefficients, and then updating the transformation parameters in an iterative fashion.

Although the above method affords greater flexibility to capture unpredictable spatial variations in the metabolite profiles, there may be cases where an even greater degree of data autonomy is desired, thereby suggesting that compartment-based parametrizations be eschewed in favor of those mechanisms whereby anatomical constraints are made more implicit. In [33], an explicit continuous model for the voxel representation of the underlying spatio-spectral distribution is proposed:

$$\rho_{hsl}(\mathbf{r}, f) = \sum_{n=1}^N \mathbf{u}(\mathbf{r}_n, f) \phi(\mathbf{r} - \mathbf{r}_n), \quad (30)$$

where ϕ is a box-shaped voxel function. The expected signal can then be expressed as:

$$s(\mathbf{k}, f) = \sum_{n=1}^N \int_{\Omega \subset \mathbb{R}^3} \mathbf{u}(\mathbf{r}_n, f) \phi(\mathbf{r} - \mathbf{r}_n) e^{-2\pi j \mathbf{k} \cdot \mathbf{r}} d\mathbf{r} \quad (31)$$

$$= \sum_{n=1}^N \mathbf{u}(\mathbf{r}_n, f) \Phi(\mathbf{k}) e^{-2\pi j \mathbf{k} \cdot \mathbf{r}_n}, \quad (32)$$

where Φ is the Fourier transform of ϕ . Following correction for B_0 effects, the reconstruction is formulated as the solution to the following regularized problem:

$$\hat{\mathbf{u}}[f_l] = \arg \min_{\mathbf{u}[f_l]} \|\tilde{\mathbf{s}}[f_l] - \mathbf{s}[f_l]\|_{\ell_2}^2 + \lambda \|\mathbf{L}\mathbf{u}[f_l]\|_{\ell_2}^2. \quad (33)$$

In this case, the regularization term consists of a spatial smoothing operator, \mathbf{L} , which penalizes local intensity variations between neighboring voxels using anatomically-derived weighting factors:

$$\|\mathbf{L}\mathbf{u}\|_{\ell_2}^2 \triangleq \sum_{i=1}^N \sum_{\substack{j \in \Omega_i \\ j \neq i}} w_{ij} |u[\mathbf{r}_i] - u[\mathbf{r}_j]|^2, \quad (34)$$

where the frequency dependence has been dropped for ease of notation. In (34), Ω_i denotes the set of voxels that are spatially adjacent to \mathbf{r}_i . The weighting factors, w_{ij} , are generally pre-computed, and can be adjusted to reflect confidences in the ascertained boundary information (e.g., see Fig. 3(b)). One clear advantage of this formulation is that it allows users to directly control (via λ) the influence of the anatomically-derived prior information. Moreover, inconsistencies between structural and spectral content will tend to manifest as degraded or biased denoising performance, rather than as additional artifacts.

Until this point, constraints imposed by the methods described above have been applied solely in the spatial domain, allowing the estimated spectra to remain maximally consistent with the observed measurements, given the appointed data model. As has been discussed, this may result in unexpected spectral behavior when discordancies exist between the model and the acquired data. A number of methods have therefore sought to further constrain the reconstruction by prescribing models for both the spatial and spectral components. In [34], the object model given by (30) is supplemented by a set of box-shaped temporal frequency basis functions, ψ :

$$\rho_{ejm}(\mathbf{r}, f) = \sum_{n=1}^N \sum_{l=1}^L v(\mathbf{r}_n, f_l) \phi(\mathbf{r} - \mathbf{r}_n) \psi(f - f_l), \quad (35)$$

where ϕ is equivalent to that in (30). The predicted signal can then be expressed as:

$$s(\mathbf{k}, t) = \Phi(\mathbf{k}) \Psi(t) \sum_{n=1}^N \sum_{l=1}^L v(\mathbf{r}_n, f_l) e^{-\alpha(\mathbf{r}_n)t} e^{-2\pi j(\mathbf{k} \cdot \mathbf{r}_n + f_l t)}, \quad (36)$$

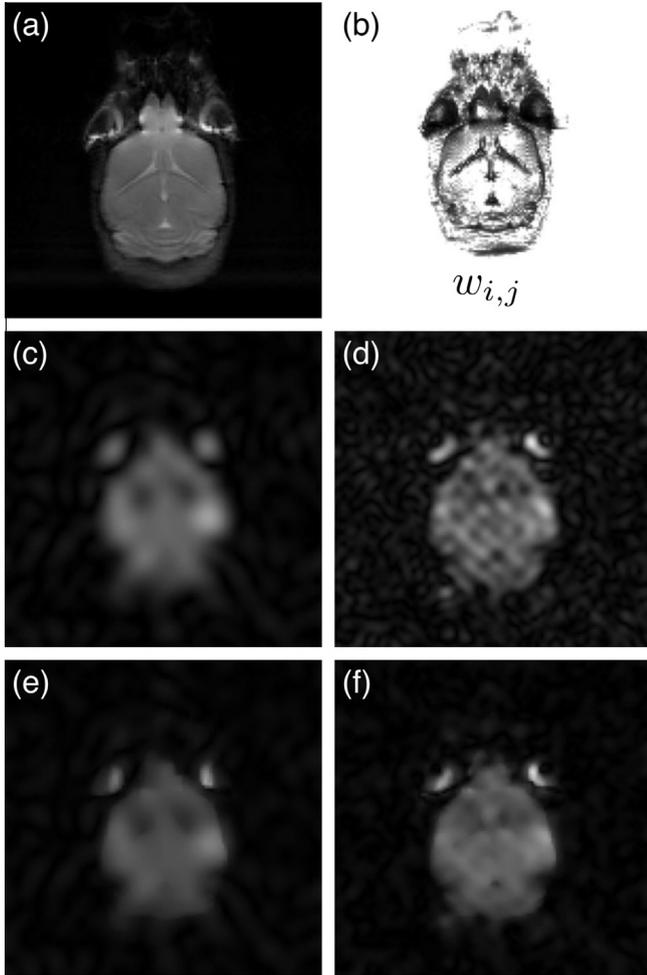


Fig. 3. Sample reconstructions of the NAA distribution of a healthy mouse brain using the procedure given by (33). (a) Anatomical image used to derive the smoothing weights, w_{ij} (b) used in the quadratic regularization penalty (34). Here, higher intensities imply stronger spatial smoothing. (c) ZDFT reconstructions from 16×16 and (d) 32×32 phase encoding steps. (e) Constrained reconstructions using the anatomical priors from 16×16 and (f) 32×32 phase encoding steps. [Figure adapted from [33] with permissions from John Wiley and Sons, Inc.]

where,

$$\alpha(\mathbf{r}_n) = \frac{1}{T_2^*(\mathbf{r}_n)} + 2\pi j\bar{\beta}(\mathbf{r}_n), \quad (37)$$

is estimated from ancillary high-resolution magnitude and phase images prior to the reconstruction, with $\bar{\beta}(\mathbf{r}_n)$ representing the mean frequency shift in voxel \mathbf{r}_n . Analogous with Φ in (32), Ψ represents the Fourier transform of the frequency domain voxel function, ψ . Given the anticipated MR spectral profile as consisting of Lorentzian lineshapes as well as smoothly varying baseline components (e.g., residual water, lipids, short T_2 macromolecules, etc.), the spectrum in each voxel is modeled as a superposition of Dirac delta functions (δ) and Chebyshev polynomials (T_n):

$$v(\mathbf{r}_n, f_l) = \underbrace{\sum_{p=1}^L w_p(\mathbf{r}_n) \delta(f_l - p)}_{\text{metabolite peaks (Dirac delta functions)}} + \underbrace{\sum_{p=1}^S W_{L+p}(\mathbf{r}_n) T_p(\mathbf{r}_n)}_{\text{broad baseline components (Chebyshev polynomials)}}. \quad (38)$$

The reconstruction is then formulated as:

$$\hat{\mathbf{w}}[f_l] = \arg \min_{\mathbf{w}[f_l]} \|\tilde{\mathbf{s}}[f_l] - \mathbf{s}[f_l]\|_{\ell_2}^2 + \lambda_1 \text{TV}_\Omega(\mathbf{w}[f_l]) + \lambda_2 \|\mathbf{w}[f_l]\|_{\ell_1}, \quad (39)$$

Sample reconstructions are provided in Fig. 4. In addition to the familiar ℓ_1 penalty, which regulates the sparsity of the weighting coefficients in (38), (39) makes use of another regularizer often espoused in the signal and image processing communities known as the total variation (TV) semi-norm, which was introduced in [35] as a means of measuring first degree information content in an input signal. TV is generally defined as:

$$\text{TV}_\Omega(f) \triangleq \sup \left\{ \int_\Omega f(\mathbf{r}) \operatorname{div} \mathbf{v}(\mathbf{r}) \, d\mathbf{r} : \mathbf{v} \in C_c^1(\Omega, \mathbb{R}^n), \|\mathbf{v}\|_{L^\infty} \leq 1 \right\}, \quad (40)$$

where $C_c^1(\Omega, \mathbb{R}^n)$ is the set of continuously-differentiable vector-valued functions compactly supported in Ω . However, if $f \in C_c^1(\Omega, \mathbb{R})$, then TV can be equivalently expressed as [36]:

$$\text{TV}_\Omega(f) = \int_\Omega |\nabla f(\mathbf{r})| \, d\mathbf{r}. \quad (41)$$

In discrete settings, for a multidimensional signal subscripted by d indices, $\mathbf{x} \in \mathbb{R}^{N_1 \times \dots \times N_d}$, TV is typically expressed as one of either:

$$\text{TV}_\Omega^{\text{iso}}(\mathbf{x}) = \sum_{i_1, i_2, \dots, i_d \in \Omega} \left| \left(\sum_{l=1}^d |D_l \{\mathbf{x}\}|^2 \right)^{1/2} \right| \quad (\text{isotropic TV}), \quad (42)$$

or,

$$\text{TV}_\Omega^{\text{ani}}(\mathbf{x}) = \sum_{i_1, i_2, \dots, i_d \in \Omega} \left| \sum_{l=1}^d |D_l \{\mathbf{x}\}| \right| \quad (\text{anisotropic TV}), \quad (43)$$

where D_l represents a discrete difference operator along the l th dimension, such that:

$$[D_l \{\mathbf{x}\}]_{i_1, i_2, \dots, i_d} = [\mathbf{x}]_{i_1, i_2, \dots, i_d} - [\mathbf{x}]_{i_1, i_2, \dots, i_d - 1, \dots, i_d}. \quad (44)$$

Much of the allure of TV rests in its proclivity for removing noise-like artifacts while preserving discerning features of an input signal, and has gained notoriety due to its versatility in addressing a wide array of problems such as denoising [37–39], suppression of truncation artifacts/ deblurring [40,41], and inpainting for sensitivity maps [42]. The reconstruction given by (39) actually introduces a slight modification of the classical TV penalty, allowing for the incorporation of tissue boundary information derived from accompanying high-resolution structural images. This is accomplished by partitioning the accompanying structural scans into K pairwise disjoint regions that are each assumed to contain spatially smooth spectral profiles, such that $\Omega = \bigcup_K \Omega_k$, thereby leading to the definition: $\text{TV}_\Omega \triangleq \sum_{k=1}^K \text{TV}_{\Omega_k}$.

Incidentally, the K-Bayes method for MRSI reconstruction [43] was proposed nearly concurrently with [34], stipulating a similar voxel-based signal model to (27). However, rather than a regularized optimization problem, the reconstruction is formulated within a Bayesian framework. In this case, the acquisition model is considered as a likelihood function, with pre-computed estimates of relaxation and field inhomogeneity parameters, while the expected spatial distribution of spectral content within anatomically-defined tissue boundaries are assimilated into a prior distribution through Markov Random Field (MRF) models. The reconstruction then consists of maximizing the posterior distribution via an expectation-maximization (EM) approach.

As mentioned previously, one of the benefits of these types of model-based and variational reconstruction frameworks is that they remain, by and large, independent of the elected k -space trajectory. This flexibility facilitates partnerships with the wealth of accelerated acquisition schemes that have been developed alongside new reconstruction approaches, including those predicated on parallel imaging techniques such as SENSE and GRAPPA [44], where multichannel information is exploited in order to correct for aliasing artifacts incurred through systematic undersampling.

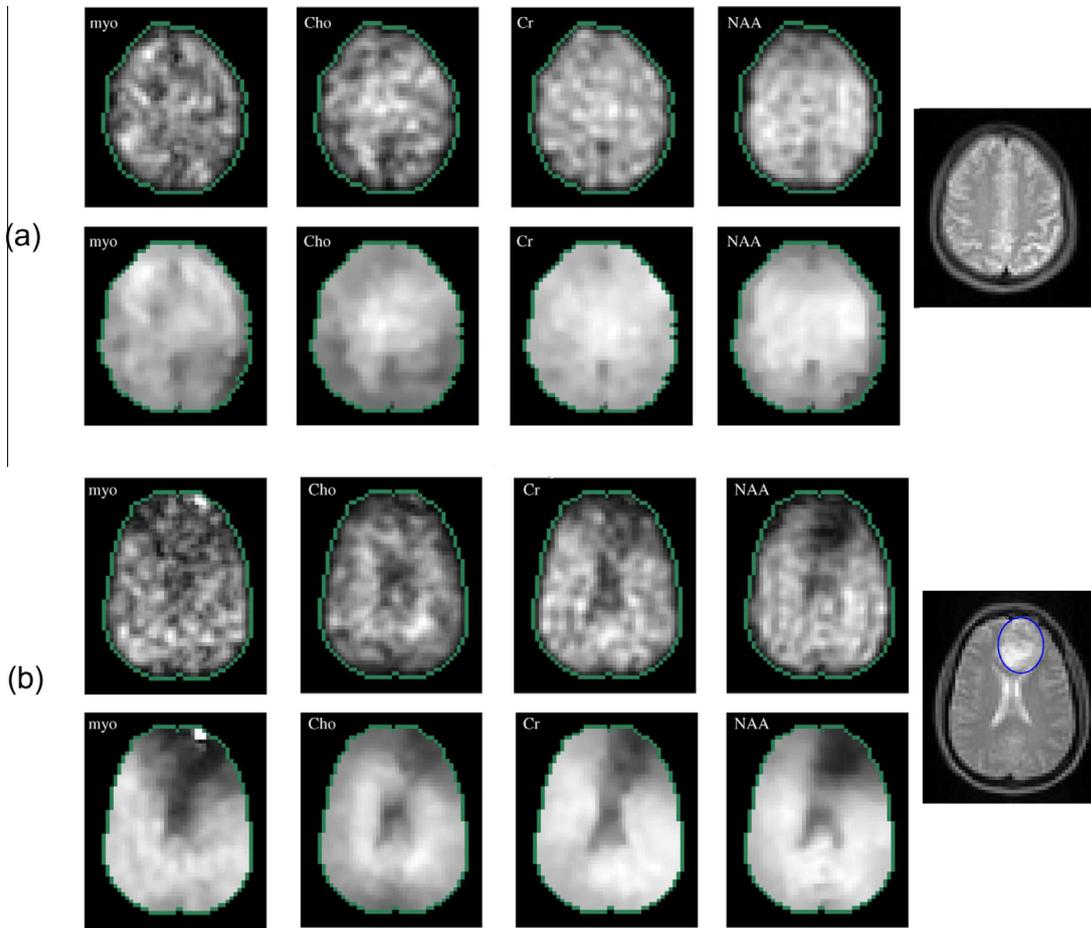


Fig. 4. Sample reconstructed metabolite maps of NAA, creatine (Cr), Cho, and myo-inositol (myo) using the method of [34] for (a) a healthy control, and (b) a clinical case following resection of a brain tumor. In each case, images in the top row represent the standard ZDFT reconstruction scheme along with post-processing lipid extraction, peak alignment, and temporal apodization, whereas the bottom rows highlight reconstruction results using the proposed method. The green lines highlight a region of interest (ROI) used to delimit *a priori* metabolite-containing areas. Attendant structural reference scans for each case are displayed in the rightmost column. For the clinical case, the approximate resection area is indicated by the blue circular region. [Figure adapted from [34] with permissions from IEEE.]

In the particular case of SENSE, where a number of studies have capitalized on the additional high-resolution information furnished through coil sensitivity maps, the predicted MRSI signal is given by:

$$s_{\zeta}(\mathbf{k}, t) = \int_{-\infty}^{\infty} \int_{\Omega \subset \mathbb{R}^3} \zeta_{\zeta}(\mathbf{r}) \rho(\mathbf{r}, f) e^{-2\pi j(\mathbf{k} \cdot \mathbf{r} + ft)} d\mathbf{r} df, \quad (45)$$

where ζ_{ζ} denotes the spatial sensitivity profile associated with the ζ^{th} coil. A fully discretized model for the observed measurements is similarly given by:

$$\tilde{\mathbf{s}}[t_i] = \mathbf{E}\mathbf{c}[t_i] + \mathbf{n}[t_i], \quad (46)$$

where the terms pertaining to the individual coil elements are concatenated along the rows of $\tilde{\mathbf{s}}$, \mathbf{E} , and \mathbf{n} . In the original weak SENSE formulation [17], compensation for aliasing artifacts is performed at the center of each reconstructed voxel. While this procedure may suffice for structural MRI resolutions, variations in the sensitivity profiles over typical MRSI voxel sizes may preclude proper unfolding, thereby giving rise to additional artifacts. This potential limitation was addressed in [45] by over-discretizing the encoding operator (\mathbf{E} in (46)), and then solving the resulting LS problem, thereby granting full access to the high-resolution sensitivity maps. Alternatively, in [46], dubbed “superresolution SENSE” (SURE-SENSE), proton echo-planar spectroscopic imaging (PEPSI) [47] data was collected from central k -space regions at the Nyquist rate, as

well as high-resolution coil sensitivity maps. The reconstruction was then given by:

$$\hat{\mathbf{c}}[t_i] = \arg \min_{\mathbf{c}[t_i]} \|\tilde{\mathbf{s}}[t_i] - \mathbf{E}\mathbf{c}[t_i]\|_{\ell_2}^2 + \lambda \|\mathbf{c}[t_i]\|_{\ell_2}^2, \quad (47)$$

where $\hat{\mathbf{c}}[t_i]$ is estimated on the same spatial grid as the sensitivity maps (for example reconstruction results, see Fig. 5).

Although the above methods address the limited unfolding capacity or resolution constraints, they do not directly confront the lingering truncation artifacts that may undermine the allure of sensitivity encoded spectroscopic imaging. Integrating the over-discretization approach of [45], the method of [48] ventures to directly control the SRF by minimizing deviations from a voxel-specific target function, \mathbf{T} (e.g., Gaussian, Dirac, etc.). Here, rather than invoking the LS framework to directly estimate the metabolite maps, the encoding scheme (or rather, its adjoint) itself is estimated:

$$\hat{\mathbf{F}} = \arg \min_{\mathbf{F}} \|\mathbf{F}\mathbf{E} - \mathbf{T}\|_F^2 + \lambda (\mathbf{F}\Psi\mathbf{F}^H), \quad (48)$$

where Ψ is the coil noise covariance matrix of the system. Under this formulation, the regularization parameter, λ , regulates the tradeoff between conformity with the designated target function and SNR performance. By choosing target functions that are fairly compact, the leakage artifacts normally attributable to significant SRF side lobes can be significantly attenuated. Once the SRF has

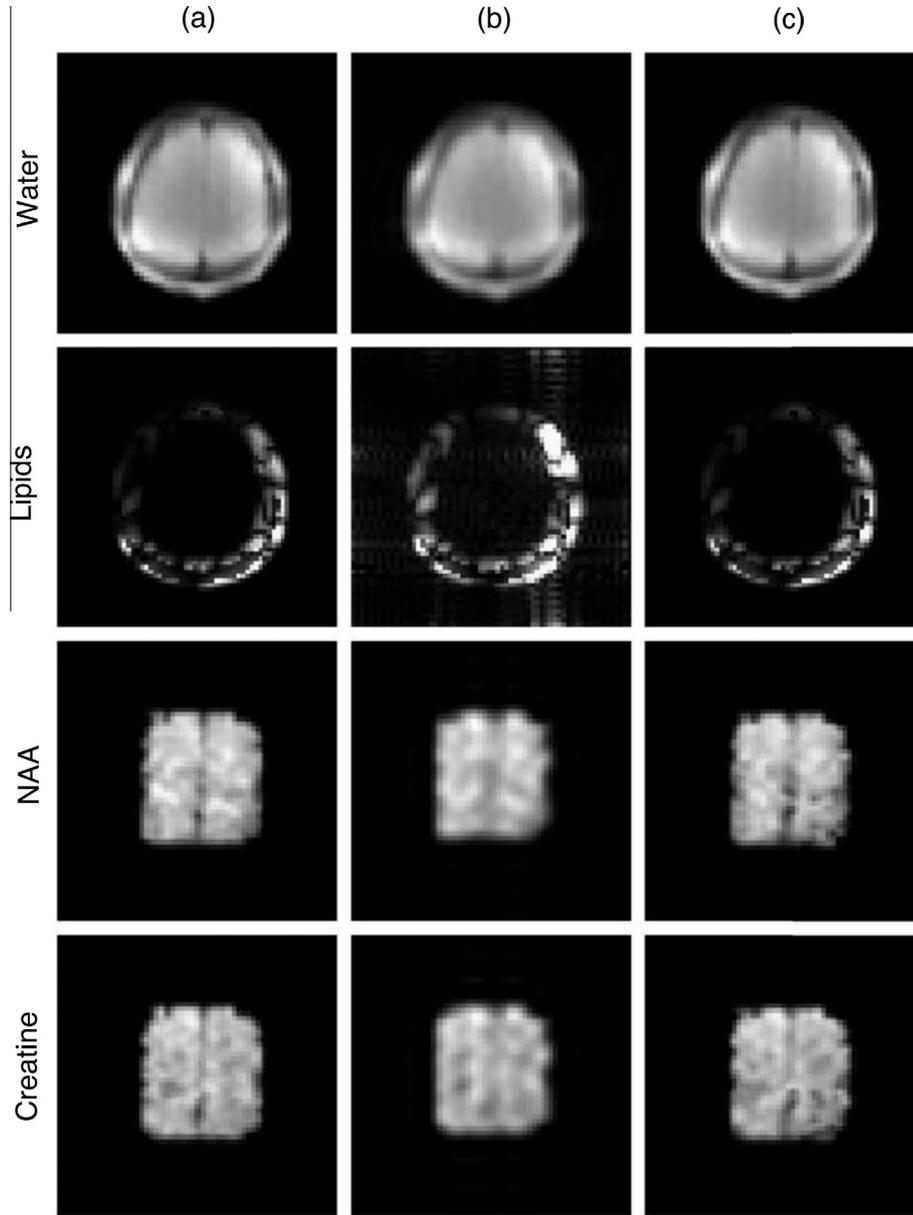


Fig. 5. Sample performance highlights of the SURE-SENSE method [46]. Reconstructions were performed by (a) inverse DFT of fully sampled data (64×64), (b) ZDFT following $2 \times$ spatial undersampling, and (c) SURE-SENSE using the $2 \times$ undersampled data. Water and lipid images were generated via full spectral integration, while the NAA and Cr maps were obtained following spectral fitting procedures. [Figure adapted from [46] with permissions from Elsevier.]

been suitably optimized, the final reconstruction is achieved by applying the estimated operator to the measured data:

$$\hat{\mathbf{c}}[t_i] = \hat{\mathbf{F}}\tilde{\mathbf{s}}[t_i]. \quad (49)$$

Thus far, each of the described methods have performed the reconstruction sequentially, separately estimating spatial maps corresponding to each individual sample in the temporal (or temporal frequency) domain. Given the often large storage requirements for MRSI datasets, this approach is typically memory efficient, yet inevitably disregards any inherent spatio-temporal correlations by implicitly asserting that each timepoint is independent. As hardware and computational power have improved, reconstruction methodologies have tended towards more multivariate approaches, seeking to exploit the full k - f volumetric information.

In [49], a variable density spiral acquisition scheme was employed to generate both the MRSI data, as well as a high-resolution water reference dataset. The latter was used to

derive estimates of the field inhomogeneity profile and coil sensitivity maps, which were then subsumed by the encoding scheme, \mathbf{E} , as well as spatio-spectral support regions (brain metabolites and extra-cranial lipids in the published manuscript). The full reconstruction was formulated as:

$$\hat{\mathbf{C}} = \arg \min_{\mathbf{C}} \|\tilde{\mathbf{S}} - \mathbf{E}\mathbf{C}\|_F^2 + \lambda_1 \underbrace{\left\| \sum_{l=1}^d \mathbf{D}_l \mathbf{C}^2 \right\|_S}_{\text{spatial smoothness}} + \lambda_2 \underbrace{\|\mathbf{C}\|_S}_{\text{sparsity}}, \quad (50)$$

where the two regularization penalties were restricted to the metabolite support region. Here, \mathbf{D}_l denotes the matrix representation of the discrete differences operator given by (44). For completeness, $\tilde{\mathbf{S}} = (\tilde{\mathbf{s}}[f_1], \tilde{\mathbf{s}}[f_2], \dots, \tilde{\mathbf{s}}[f_L])$, and $\mathbf{C} = (\mathbf{c}[f_1], \mathbf{c}[f_2], \dots, \mathbf{c}[f_L])$. In this case, the first regularizer (controlled via λ_1) compels spatial smoothness by promoting spatial maps with sparse gradients,

whereas the second (controlled via λ_2) aims to limit any noise precipitated during the iterative reconstruction procedure.

6.1. From voxel to feature-based interpretations

One of the advantages proffered by working directly with the full MRSI dataset is that highly correlated measurements can often be distilled down to a modest number of prominent spatio-spectral features. Indeed, some of the earliest endeavors in non-traditional MRSI reconstruction such as SLIM ultimately subscribe to this very notion, often attempting to improve reconstruction quality or elicit further gains in other aspects of the measurement process by minimizing data redundancy (e.g., SLOOP). While it may be argued that the SLIM formulation itself often remains too restricted to fully accommodate the heterogeneities underlying MRSI measurements, its simplicity and fundamental object characterization have long maintained their appeal. As testament, a growing number of contemporary methods have espoused models of the form:

$$\rho_{\text{bilin}}(\mathbf{r}, f) = \sum_{\kappa=1}^K \mathbf{u}_{\kappa}(\mathbf{r}) v_{\kappa}(f). \quad (51)$$

Also referred to as the “partially separable” model [50], expression (51) posits a bilinear decomposition of the spatio-spectral distribution reminiscent of SLIM-type methods, effectively decoupling spatial and temporal behavior. However, contrary to SLIM, where the objective is to estimate the unknown spectral components (v_{κ}) assuming that the spatial components (u_{κ}) are known *a priori*, the focus has shifted to finding general decompositions in the form of (51) that yield low model orders (i.e., small K). This approach can

be justified for MRSI applications by considering the following limiting cases. In one scenario, each v_{κ} could correspond to an individual resonance peak, with u_{κ} its accompanying spatial distribution. In this case, the model order will be determined by the number of distinguishable spectral peaks, K_r , which typically only number up to around 20–30 in common MRSI experimental settings, depending on the measured nucleus, echo time selection, etc. Alternatively, the object may be fully described by only K_f unique features, such as the spectrally homogeneous compartments in the SLIM model, or spectral signatures associated with commonly observed metabolites. Hence, for a given object, the underlying model order can be expected to reside somewhere in $[K_f, K_r]$.

Given sufficient sampling conditions, the model order (similarly labeled as the data “dimensionality”) can often be inferred from the measured data through the use of so-called “rank-revealing” transforms, such as the singular value decomposition (SVD), by examining the resulting eigenvalue spectra. The MR community has recently witnessed an outburst of reconstruction methodologies seeking to capitalize on the surmised low-dimensional structure underlying seemingly complex objects and processes by specifically promoting low-rank solutions, mostly notably in the field of dynamic MRI, e.g., [51–56]. Although these methods certainly each carry their own distinctions, approaches have typically fallen into one of two categories: (I) The reconstruction is performed over the composite spatio-temporal dataset, whereby low-rank solutions are exacted explicitly, or through the use of additional regularizing penalties. (II) The temporal basis functions are first estimated from limited k -space data acquired at high temporal resolution (e.g., using the SVD), which are then used to guide

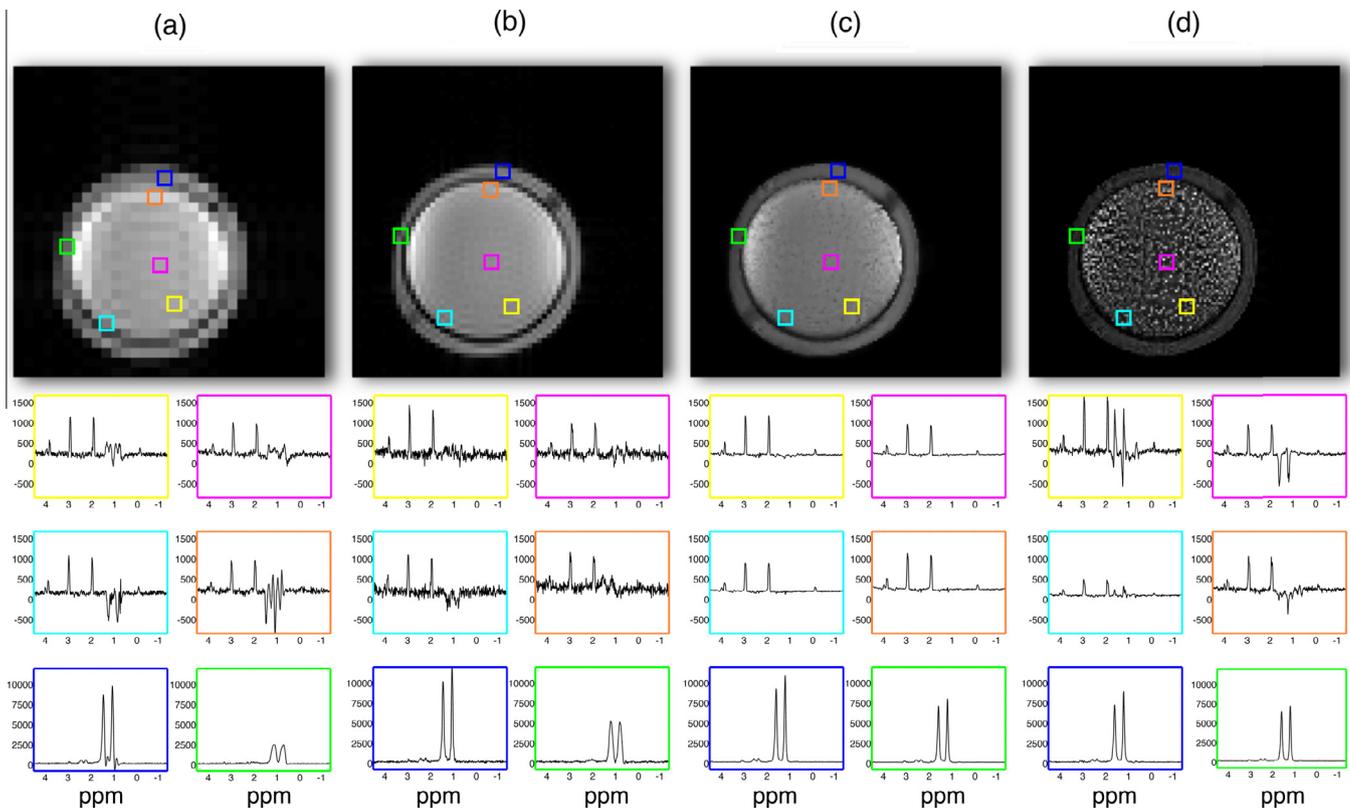


Fig. 6. Reconstruction results on a phantom using the method of [59]. The phantom consisted of two compartments: an inner spherical compartment contained a solution of 50 mmol/L NAA and 50 mmol Cr in doped water. The outer cylindrical compartment was filled with corn oil. (a) Standard DFT reconstruction from 32×32 and (b) 64×64 CSI encoding steps. (c) Reconstruction using the proposed technique. Here, the efficacy of the spatial regularization can be clearly seen through comparison with (d), which demonstrates reconstruction outcome when λ is set to zero in (58). Magnitude images were obtained via full spectral integration. Representative spectra are also provided, corresponding to the spatial locations denoted by the colored boxes in the spatial images. [Figure adapted from [59] with permissions from IEEE.]

the recovery of attendant spatial basis functions from either a separate, or combined acquisition with greater k -space coverage but limited temporal sampling.

Despite the reported success of these approaches, a number of crucial disparities confound straightforward translation to MRSI settings. Firstly, low metabolite SNR, as well as the decaying exponential envelopes that characterize spectroscopic signals render rank-based interpolation schemes particularly susceptible to noise-related artifacts and biases. Secondly, unless explicitly accounted for, field inhomogeneity or off-resonance effects present in the estimated temporal basis functions may bias the recovery of the corresponding spatial components. Nonetheless, approach (I) has been explored as a part of a general denoising scheme for CSI acquisitions labeled LORA [57], which can be used to improve SNR in datasets where sufficient signal averaging is relegated in favor of measurement time or increased k -space coverage. First, the data are filtered by finding the nearest (in an ℓ_2 sense) rank- K approximation to the acquired measurements:

$$\hat{\mathbf{S}} = \arg \min_{\mathbf{S}} \|\tilde{\mathbf{S}} - \mathbf{S}\|_F^2, \quad \text{s.t. rank}(\mathbf{S}) = K. \quad (52)$$

Expression (52) can be solved by first taking the singular value decomposition of the measurement matrix:

$$\tilde{\mathbf{S}} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^H, \quad (53)$$

where for $\tilde{\mathbf{S}} \in \mathbb{C}^{(M \times T)}$, $\mathbf{U} = (\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_M) \in \mathbb{C}^{(M \times M)}$, $\mathbf{V} = (\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_T) \in \mathbb{C}^{(T \times T)}$, are unitary matrices containing the left- and right-singular vectors of $\tilde{\mathbf{S}}$, respectively, and $\mathbf{\Sigma} \in \mathbb{R}_+^{(M \times T)}$ contains the corresponding singular values, σ_i along the diagonal. Solutions to (52) can then be found by approximating $\tilde{\mathbf{S}}$ using only the first K singular vectors/values:

$$\hat{\mathbf{S}} = \sum_{\kappa=1}^{K \leq \min(M,T)} \sigma_{\kappa} \mathbf{u}_{\kappa} \mathbf{v}_{\kappa}^H. \quad (54)$$

Following the low-rank approximation scheme (which can be considered as a global denoising step), the data are then Fourier transformed back into the spatial domain, and additional denoising is performed on each of the voxel spectra individually. Here, the spectra are assumed to consist of at most P Lorentzian peaks, in principle leading to a discrete time-domain signal that is P th-order linearly predictable, such that:

$$\hat{s}[t_i] = \sum_{p=1}^P \zeta_p s[t_i - t_p], \quad (55)$$

where the spatial variable has been omitted for ease of presentation. A necessary and sufficient condition for (55) is that the Hankel matrix formed from its samples:

$$\mathbf{H}' = \begin{pmatrix} \hat{s}[t_1] & \hat{s}[t_2] & \cdots & \hat{s}[t_L] \\ \hat{s}[t_2] & \hat{s}[t_3] & \cdots & \hat{s}[t_{L+1}] \\ \vdots & \vdots & \ddots & \vdots \\ \hat{s}[t_{T-L+1}] & \hat{s}[t_{T-L+2}] & \cdots & \hat{s}[t_T] \end{pmatrix}, \quad (56)$$

for some $P < L < T$, should be rank K [58]. Additional denoising of the individual voxel spectra then amounts to solving:

$$\hat{\mathbf{H}} = \arg \min_{\mathbf{H}} \|\mathbf{H}' - \mathbf{H}\|_F^2, \quad \text{s.t. rank}(\mathbf{H}) = P, \quad (57)$$

for each row in $\hat{\mathbf{S}}$ using the same procedure outlined by (53) and (54). The denoised spectrum is then obtained by extracting the first row and last column of $\hat{\mathbf{H}}$.

Aside from providing a “natural” infrastructure for the spatio-spectral distribution, the decoupling between spatial and temporal characteristics in the bilinear model allows assumptions

surrounding signal characteristics in either domain to be selectively applied. For example, with the aid of a separately acquired high-resolution static field inhomogeneity map, [59] encourages the recovery of spatial components that are expressible as real-valued, non-negative, piecewise linear functions. This is partially achieved through the use of a total generalized variation penalty [60,37], which extends the classical TV penalty (40) and (41) to accommodate higher-order derivative information. Indeed, the use of higher-order functionals in deblurring and denoising applications (for additional examples, see [61–64]) has been extensively researched due to their ability to evade the so-called “staircasing” artifacts often encountered when considering only first order information [65]. In order to ensure access to the individual spatial and spectral components, the reconstruction is formulated as the joint estimation problem:

$$\hat{\mathbf{U}}, \hat{\mathbf{V}} = \arg \min_{\mathbf{U}, \mathbf{V}} \|\tilde{\mathbf{S}} - \mathcal{FB}\{\mathbf{UV}\}\|_F^2 + \lambda \text{TGV}_{\alpha}^2(\mathbf{U}) \quad \text{s.t. } \mathbf{U} \in \mathcal{C}_U, \mathbf{V} \in \mathcal{C}_V. \quad (58)$$

where \mathcal{F} and \mathcal{B} are spatio-temporal operators representing the encoding scheme and the spatially-dependent phase perturbations introduced by the static field inhomogeneity, respectively. Similarly, \mathbf{U} and \mathbf{V} are discrete matrix representations of u_{κ} and v_{κ} in (51), respectively, TGV_{α}^2 is the total generalized variation of second order, and \mathcal{C}_U and \mathcal{C}_V are constraint sets – the significance of which shall be discussed shortly. The key to the method is that the separability of (51) (i.e., of \mathbf{U} and \mathbf{V}) is strictly enforced during the reconstruction. This is necessary due to the fact that the presence of field inhomogeneities may induce coupling between the spatial and spectral components, which typically results in \mathbf{U} subsuming an additional time dependence. Therefore, in order to effectively segregate the two domains, and to comply with the assumptions that the spatial components are real-valued and non-negative, the current estimate of \mathbf{U} is projected at each iteration of the reconstruction onto the set \mathcal{C}_U , which encompasses two sub-operations such that:

$$\mathcal{P}_{\mathcal{C}_U}(\mathbf{U}) = \mathcal{P}_{\mathbb{R}^+} \mathcal{P}_{\Sigma}(\mathbf{U}), \quad (59)$$

where,

$$\mathcal{P}_{\mathbb{R}^+}(\mathbf{X}) = \max\left(0, \frac{1}{2}(\mathbf{X} + \mathbf{X}^*)\right), \quad (60)$$

denotes the projection onto the real-valued, non-negative orthant and,

$$\mathcal{P}_{\Sigma}(\mathbf{X}) = \frac{1}{T} \sum_{i=1}^T \mathbf{X}[t_i], \quad (61)$$

a projection onto the mean temporal timecourse. The temporal components themselves are similarly restricted to the unit Frobenius norm by projecting onto the constraint set \mathcal{C}_V :

$$\mathcal{P}_{\mathcal{C}_V}(\mathbf{V}) = \begin{cases} \frac{\mathbf{V}}{\|\mathbf{V}\|_F}, & \|\mathbf{V}\|_F \geq 1 \\ \mathbf{V}, & \text{otherwise,} \end{cases} \quad (62)$$

which has been shown to minimize scaling ambiguities during the reconstruction [66–68]. Sample reconstruction results on a two-compartment MR phantom can be viewed in Fig. 6.

While independence of acquisition strategy is a laudable trait that affords flexibility to a reconstruction routine, methods subscribing to approach (II) intentionally build interdependencies in order to establish a framework that captures the most desirable characteristics during each stage of the process. As an example, the Spectroscopic Imaging by Exploiting Spatio-Spectral Correlation (SPICE) method [69], aims to enhance the SNR of accelerated and extended k -space acquisition schemes by leveraging the denoising capabilities of low-rank modeling, and consists of two

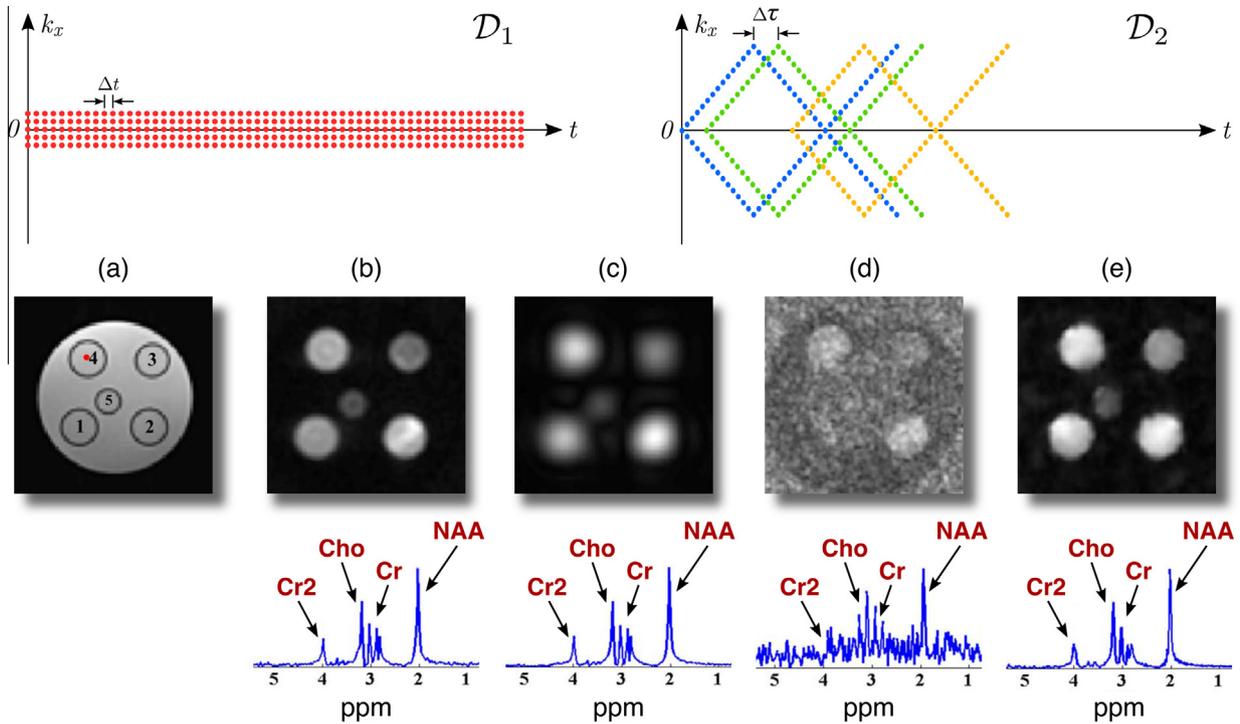


Fig. 7. Overview of the SPICE acquisition protocol and sample reconstruction results. (Top row), example acquisition strategies for acquiring \mathcal{D}_1 (top left) and \mathcal{D}_2 (top right). Middle row: (a) structural reference image demonstrating the phantom geometry, which consisted of five vials placed in a cylindrical jar filled with doped water. The contents of the vials were as follows: 1 (20 mmol/L NAA, 15 mmol/L Cr, 10 mmol/L Cho), 2 (20 mmol/L NAA, 10 mmol/L Cr, 5 mmol/L Cho), 3 (10 mmol/L NAA, Cr, Cho, myo), 4 (15 mmol/L NAA, 8 mmol/L Cr, 5 mmol/L Cho), with vial 5 containing the same solution as vial 3. The NAA maps shown were reconstructed using (b) DFT from CSI data using 60×60 encoding steps, (c) ZDFT from CSI data using 19×19 encodings, (d) EPSI data using 100×100 encoding steps and two averages, (e) the SPICE method using 12×12 CSI encodings to generate \mathcal{D}_1 , and 45 shifted echoes for \mathcal{D}_2 . Representative spectra originating from vial 4 (denoted by the small red dot in (a)) corresponding to each reconstruction method are illustrated in the bottom row. [Figure adapted from [69] with permissions from John Wiley and Sons, Inc.]

main modules. First, a MRSI dataset \mathcal{D}_1 is obtained with traditional Cartesian sampling, acquiring only a small central region in k -space but with high temporal resolution. This dataset is then corrected for field inhomogeneities, whereupon a set of K temporal basis functions, $\hat{\mathbf{V}}$, are estimated using (53). A second dataset \mathcal{D}_2 is then collected with extended k -space coverage but reduced temporal sampling, for example via accelerated strategies such as spiral [70–72] or EPSI [47]. The corresponding spatial basis functions are then estimated by solving:

$$\hat{\mathbf{U}} = \arg \min_{\mathbf{U}} \|\hat{\mathbf{S}}_{\mathcal{D}_2} - \mathcal{FB}\{\mathbf{U}\hat{\mathbf{V}}\}\|_F^2 + \lambda \|\mathcal{L}\{\mathbf{U}\hat{\mathbf{V}}\}\|_F^2. \quad (63)$$

Here, $\hat{\mathbf{S}}_{\mathcal{D}_2}$ are the measurements corresponding to dataset \mathcal{D}_2 , and \mathcal{L} is an anatomically-derived smoothing penalty analogous to \mathbf{L} in (34). \mathcal{F} and \mathcal{B} are defined equivalently as in (58). Note that unlike previously described parametric methods that estimate one set of basis functions from accompanying extra-modal scans, the “*a priori*” information utilized with these type (II) methods is derived in a data consistent manner from intra-modal acquisitions, thereby reducing the potential for model mis-match artifacts. A schematic representation of various acquisition schemes that may be adopted to acquire \mathcal{D}_1 and \mathcal{D}_2 , along with sample reconstructions can be found in Fig. 7. The reconstruction given by (63) was also later extended to include non-quadratic regularization penalties, which was reported to lead to improved reconstruction performance [73].

The SPICE method nicely illustrates the potential benefits that can be attained through a judicious pairing of the acquisition and reconstruction, and highlights an increasing trend in the MR community. Other such examples can be found in the compressed sensing (CS) literature, where alternative acquisition strategies are required in order to comply with the underlying theoretical tenets. While CS has typically been regarded primarily as a means

for further reducing k -space sampling requirements (e.g., [74–76]), it has also shown promise as a potential avenue for spatial resolution enhancement. In [77], for example, a 3D flyback MRSI sequence was utilized for hyperpolarized ^{13}C applications, whereupon spectral domain sparsity was exploited during reconstruction in order to achieve a twofold increase in spatial resolution.

7. Practical limitations

While the methods outlined above demonstrate a clear capacity for improved reconstruction quality, it is important to acknowledge any prospective shortcomings or limitations. For example, with any model-based approach, reconstruction outcome will ultimately be circumscribed by the accuracy of the underlying model. As such, failure to accommodate potentially obtrusive experimental confounds such as field inhomogeneity, eddy currents, patient motion, or non-uniform rf profiles may bias the reconstructions in an unpredictable or inscrutable manner. Furthermore, any discordancies between the measured MRSI data and ancillary extra-modal *a priori* information, for example those due to segmentation or registration inaccuracies, may manifest as additional artifacts or structured noise patterns that may be difficult to distinguish. Another potential drawback for non-linear reconstruction methods is that most of the standard tools commonly employed for the analysis of linear systems, such as direct access to the SRF, are no longer applicable or available, thereby obfuscating robust method characterization.

Although the variational/inverse problem formulation is able to successfully bypass many of the limitations prescribed by the inverse DFT reconstructions themselves, most of the described methods rely upon similar data pre-processing measures, such as

residual water and/or lipid suppression in ^1H applications, in order to achieve reasonable results. This is due to the fact that reconstruction outcomes are typically driven by the dominant signal components, which may bias the reconstruction of metabolite maps of interest if primarily attributable to large nuisance signals. Similarly, whereas the collection of an auxiliary water reference scan for eddy current and static field inhomogeneity correction is feasible at current MRSI resolutions, the additional time required for such measurements at high resolution far exceed clinically reasonable exam durations.

Finally, it should be noted that while absolute quantitation has typically been considered as a separate issue beyond the scope of the majority of the discussed works, and by extension, of the present survey, the accurate recovery and preservation of spectral features is paramount. Indeed, although many studies often eschew absolute quantification altogether, instead reporting significant alterations or aberrations in metabolite ratios, most superresolution reconstruction paradigms currently lack robust mechanisms for safeguarding such interrelationships, thereby carrying the risk of biased results. It is therefore imperative to ensure that certain resonances are not over or under-represented by the reconstruction procedure during validation stages.

8. Translation of superresolution reconstruction methods into clinical settings

Despite the profusion of promising methodologies, the clinical and clinical research communities seldom deviate from traditional inverse DFT reconstructions. While the rationale underlying such indispositions may stem from a complex and diverse array of factors, a number of key potential explanations are provided in the following.

Ultimately, despite their limitations, standard DFT (including ZDFT) reconstructions remain the most popular due to their computational ease and simplicity. In point of fact, in contrast to many of the superresolution methods described earlier, standard DFT reconstruction outcome is neither reliant upon exact knowledge of the elected acquisition scheme, nor the congruity of the measured data with externally derived prior information, thereby avoiding complications related to issues such as segmentation and registration. Moreover, DFT reconstructions require little user interaction and no special parameter tuning, perhaps one of the more arcane and daunting aspects of the regularization-based reconstruction frameworks for non-experts. Indeed, such reservations are defensible, as the regularization parameter (λ) ultimately regulates the bias-variance tradeoff, effectively determining the overall behavior of a given reconstruction method. Although a considerable body of research has been dedicated towards establishing robust criteria for selecting λ , universal consensus is thus far lacking. Nonetheless, established approaches – mostly envisioned for use with quadratic penalties – include the discrepancy principle [78,79], generalized cross validation [80], and the L-curve method [81], with more detailed discussions found in [82,83]. More recent studies have explored alternative means for selecting λ in non-linear reconstruction settings, such as those marked by their use of ℓ_1 [84] and TV penalties [85–87], as well as through more general adaptive frameworks such as Stein's unbiased risk estimate (SURE) [88–91]. In general, an appropriate choice of λ follows from careful deliberations involving the data noise variance, and the (eigen) spectral properties of the encoding operators, neither of which may be known *a priori*. Similar obstacles beleaguer selection of the model order in the aforementioned rank-minimization reconstruction schemes, where robust estimation of the underlying data dimensionality from a series of noisy measurements represents an ongoing challenge in statistics, information theory, and

machine learning. Again, while no technique can boast complete infallibility, several methods have been shown to be robust in effectively segregating the signal and noise subspaces, including the minimum description length (MDL) [92], Akaike information criterion (AIC) [93], Scree test on residual variance [94], cross-validation approaches [95–97], and the Bayesian information criterion (BIC) [98,99] – including a number of probabilistic principal component analysis (PCA) approaches [100–102]. Further details, as well as a more general discussion on the problem of model order selection can be found in works such as [103].

Another factor encumbering the advancement of superresolution reconstruction methods for MRSI is the difficulty in method validation. Unlike many structural MR modalities, where suitable gold standard measurements are available for validating new techniques, as is the case with parallel imaging or compressed sensing [104,105], MRSI offers a fundamentally unique set of information that lacks a referential counterpart. This distinctiveness portends a number of challenges when considering techniques that allege to transcend current resolution capabilities, especially given the stringency required in order to meet clinical standards. Although a few studies have reported on regional metabolite distributions in the human brain at enhanced spatial resolution (e.g., [106,107]), the general dearth of high-resolution gold standard data for *in vivo* MRSI necessarily arouses a degree of skepticism at the clinical level when viewing superresolution reconstruction outcomes. Such misgivings are also often compounded by inadequate method characterization. For example, whereas the properties and potential artifacts associated with DFT reconstructions are well-established, and therefore easily identified by a trained reviewer, such claims generally cannot be made with regards to the various idiosyncrasies defining non-traditional approaches. This deficiency may arise either due to hitherto insufficient experimentation, to impediments raised by the underlying theory, or simply to the often times discordant priorities set by the image reconstruction and clinical research communities. One common avenue for addressing difficulties in both validation and characterization is the use of phantoms, i.e., objects with known geometrical and spectral properties. While unable to fully capture the inherent complexities of living biological systems, such objects – especially those designed to emulate the spatio-spectral heterogeneities and sophistication normally encountered in *in vivo* applications (e.g., [108,109]) – may provide a reliable and reproducible means for assessing the high resolution capabilities of a given method, thereby allaying a number of key clinical concerns.

Finally, it should be mentioned that despite sophisticated modern hardware, the computational burden imposed by many of the aforementioned methods often leads to reconstruction times on the order of tens of minutes to hours. While such requirements may not present a problem for clinical research studies, they do currently preclude the use of such advanced reconstruction procedures in time-sensitive clinical applications.

9. Conclusion

MRSI has always been regarded as a challenging topic due to the considerable endeavors in both promoting favorable measurement conditions, and employing the necessary post-processing routines so as to yield practicable data. Moreover, even if sufficient data quality can be achieved, resolution limitations have thus far remained the foremost impediment to the widespread clinical utility of MRSI. Over the past few decades, research efforts have been primarily directed towards hardware improvements, including higher field strength scanners to achieve the accompanying sensitivity gains, more sensitive coil arrays, and more sophisticated shimming systems, as well as accelerated acquisition techniques

in order to maximize the $k-f$ coverage per unit time. Equally valuable to the advancement of MRSI, such techniques aim to leverage any improvements in the data SNR to ostensibly augment M (and N) in the DFT matrix (7). While this approach may further align the reconstruction conditions with those surrounding traditional structural MR, it still does not take into account any potentially beneficial prior information regarding the specifics of the acquisition, or the measured object itself. It is the command of this knowledge that forms the foundation for the plenitude of superresolution reconstruction methods that have been developed in parallel, alternatively increasing N by exploiting known geometrical properties, or proposing alternative bases that may more accurately or efficiently capture salient spatio-spectral features.

Although the acquisition and reconstruction are often dissociated, perhaps the most significant advances will be ushered in by a successful partnership between the two. Indeed, while most MR experiments are conceived in a linear fashion – considering the reconstruction only after hardware and acquisition protocols have been established – it is clear that a pipeline that is able to fully capitalize on the strengths offered by each module may lead to significantly improved outcomes. For example, improved nuisance signal suppression or improved sensitivity at the acquisition level would enable variation-based reconstruction frameworks to be fully driven by the metabolic signal of interest, whereas more robust *a priori* information may allow acquisition protocols to be relaxed in order to further reduce sampling requirements and/or exam durations.

Currently, the primary obstacle when translating superresolution methods into more applied or clinical settings is the shortfall of validation and characterization work, thereby arousing apprehensions and reservations when reviewing reconstruction results. Admittedly, such misgivings may be difficult to overcome, especially considering the uniqueness of the spatio-spectral information furnished by MRSI. Extensive testing across scanner vendors and institutions, focus on the development of fully automated reconstruction procedures as well as comprehensive and robust analytical tools for suitably characterizing non-traditional (especially non-linear) techniques, and close collaborations with the clinical community are therefore vital steps. Such endeavors, coupled with modern hardware and computational advances (e.g. parallel computing) will be needed in order to develop the necessary confidences required to propel superresolution MRSI beyond its current confines.

Acknowledgments

This work was supported in part by the Center for Biomedical Imaging (CIBM), and in part by the Swiss National Science Foundation under SNF 320030-135425.

References

- [1] J. Granot, Selected Volume Excitation Using Stimulated Echoes (VEST). Applications to spatially localized spectroscopy and imaging, *J. Magn. Reson.* (1969) 70 (3) (1986) 488–492.
- [2] J. Frahm, K.-D. Merboldt, W. Hnicke, Localized proton spectroscopy using stimulated echoes, *J. Magn. Reson.* (1969) 72 (3) (1987) 502–508.
- [3] R. Kimmich, D. Hoepfel, Volume-selective multipulse spin-echo spectroscopy, *J. Magn. Reson.* (1969) 72 (2) (1987) 379–384.
- [4] P.A. Bottomley, Selective Volume Method for Performing Localized NMR Spectroscopy, US Patent, October 1984. <<http://www.freepatentsonline.com/4480228.html>>.
- [5] P.A. Bottomley, Spatial localization in NMR spectroscopy in vivo, *Ann. New York Acad. Sci.* 508 (1) (1987) 333–348.
- [6] W.-I. Jung, Localized double spin echo proton spectroscopy. Part I: Basic concepts, *Concept. Magnet. Res.* 8 (1) (1996) 1–15.
- [7] W.-I. Jung, Localized double spin echo proton spectroscopy. Part II: Weakly coupled homonuclear spin systems, *Concept. Magnet. Res.* 8 (2) (1996) 105–118.
- [8] R. Pohmann, M. von Kienlin, A. Haase, Theoretical evaluation and comparison of fast chemical shift imaging methods, *J. Magn. Reson.* 129 (2) (1997) 145–160.
- [9] S. Posse, R. Otazo, S.R. Dager, J. Alger, MR spectroscopic imaging: principles and recent advances, *J. Magn. Reson. Imag.* 37 (6) (2013) 1301–1325.
- [10] X. Hu, D.N. Levin, P.C. Lauterbur, T. Spraggins, SLIM: spectral localization by imaging, *Magn. Reson. Med.* 8 (3) (1988) 314–322.
- [11] X. Hu, Z. Wu, SLIM revisited, *IEEE Trans. Med. Imag.* 12 (1993) 583–587.
- [12] Z. Liang, P. Lauterbur, A theoretical analysis of the SLIM technique, *J. Magn. Reson., Ser. B* 102 (1) (1993) 54–60.
- [13] Z.-P. Liang, F. Boada, R. Constable, E.M. Haacke, P. Lauterbur, M.R. Smith, Constrained reconstruction methods in MR imaging, *Rev. Magn. Reson. Med.* 4 (1992) 67–185.
- [14] M. von Kienlin, R. Mejia, Spectral localization with optimal pointspread function, *J. Magn. Reson.* 94 (2) (1991) 268–287.
- [15] S.J. Reeves, Selection of K-space samples in localized MR spectroscopy of arbitrary volumes of interest, *J. Magn. Reson. Imag.* 5 (2) (1995) 245–247.
- [16] L. An, S. Warach, J. Shen, Spectral localization by imaging using multielement receiver coils, *Magn. Reson. Med.* 66 (2011) 1–10.
- [17] K.P. Pruessmann, M. Weiger, M.B. Scheidegger, P. Boesiger, SENSE: sensitivity encoding for fast MRI, *Magn. Reson. Med.* 42 (1999) 952–962.
- [18] Z.P. Liang, P.C. Lauterbur, A generalized series approach to MR spectroscopic imaging, *IEEE Trans. Med. Imag.* 10 (2) (1991) 132–137.
- [19] I. Khalidov, D. Van De Ville, M. Jacob, F. Lazeyras, M. Unser, BSLIM: spectral localization by imaging with explicit B_0 field inhomogeneity compensation, *IEEE Trans. Med. Imag.* 26 (7) (2007) 990–1000.
- [20] J.A. Kmieciak, C.D. Gregory, Z.-P. Liang, P.C. Lauterbur, M.J. Dawson, Lactate quantitation in a gerbil brain stroke model by GSLIM of multiple-quantum-filtered signals, *J. Magn. Reson. Imag.* 9 (4) (1999) 539–543.
- [21] Z. Dong, J.-H. Hwang, Lipid signal extraction by SLIM: application to ^1H MR spectroscopic imaging of human calf muscles, *Magn. Reson. Med.* 55 (6) (2006) 1447–1453.
- [22] R. Löffler, R. Sauter, H. Kolem, A. Haase, M. von Kienlin, Localized spectroscopy from anatomically matched compartments: improved sensitivity and localization for cardiac ^31P MRS in humans, *J. Magn. Reson.* 134 (2) (1998) 287–299.
- [23] M. Meininger, W. Landschütz, M. Beer, T. Seyfarth, M. Horn, T. Pabst, A. Haase, D. Hahn, S. Neubauer, M. von Kienlin, Concentrations of human cardiac phosphorus metabolites determined by SLOOP ^31P NMR spectroscopy, *Magn. Reson. Med.* 41 (4) (1999) 657–663.
- [24] M. Beer, T. Seyfarth, J. Sandstede, W. Landschütz, C. Lipke, H. Köstler, M. von Kienlin, K. Harre, D. Hahn, S. Neubauer, Absolute concentrations of high-energy phosphate metabolites in normal, hypertrophied, and failing human myocardium measured noninvasively with ^31P -SLOOP magnetic resonance spectroscopy, *J. Am. College Cardiol.* 40 (7) (2002) 1267–1274.
- [25] S. Plevritis, A. Macovski, Spectral extrapolation of spatially bounded images [MRI application], *IEEE Trans. Med. Imag.* 14 (3) (1995) 487–497.
- [26] R. Gerchberg, Super-resolution through error energy reduction, *Opt. Acta: Int. J. Opt.* 21 (9) (1974) 709–720.
- [27] A. Papoulis, A new algorithm in spectral analysis and band-limited extrapolation, *IEEE Trans. Circ. Syst.* 22 (9) (1975) 735–742.
- [28] S.K. Plevritis, A. Macovski, MRS imaging using anatomically based k-space sampling and extrapolation, *Magn. Reson. Med.* 34 (5) (1995) 686–693.
- [29] M. Jacob, X. Zhu, A. Ebel, N. Schuff, Z.-P. Liang, Improved model-based magnetic resonance spectroscopic imaging, *IEEE Trans. Med. Imag.* 26 (10) (2007) 1305–1318.
- [30] B.K. Natarajan, Sparse approximate solutions to linear systems, *SIAM J. Comput.* 24 (2) (1995) 227–234.
- [31] S. Mallat, Z. Zhang, Matching pursuits with time-frequency dictionaries, *IEEE Trans. Signal Process.* 41 (12) (1993) 3397–3415.
- [32] Y. Pati, R. Rezaifar, P. Krishnaprasad, Orthogonal matching pursuit: recursive function approximation with applications to wavelet decomposition, in: 1993 Conference Record of The Twenty-Seventh Asilomar Conference on Signals, Systems and Computers, 1993, 1, November 1993, pp. 40–44.
- [33] J.P. Haldar, D. Hernando, S.-K. Song, Z.-P. Liang, Anatomically constrained reconstruction from noisy data, *Magn. Reson. Med.* 59 (4) (2008) 810–818.
- [34] R. Eslami, M. Jacob, Robust reconstruction of MRSI data using a sparse spectral model and high resolution MRI priors, *IEEE Trans. Med. Imag.* 29 (6) (2010) 1297–1309.
- [35] L.I. Rudin, S. Osher, E. Fatemi, Nonlinear total variation based noise removal algorithms, *Phys. D: Nonlinear Phenom.* 60 (1992) 259–268.
- [36] T. Chan, S. Osher, J. Shen, The digital TV filter and nonlinear denoising, *IEEE Trans. Image Process.* 10 (2) (2001) 231–241.
- [37] F. Knoll, K. Bredies, T. Pock, R. Stollberger, Second order Total Generalized Variation (TGV) for MRI, *Magnet. Reson. Med.* 65 (2) (2011) 480–491.
- [38] Y. Wang, H. Zhou, Total variation wavelet-based medical image denoising, *Int. J. Biomed. Imag.* (2006) 6.
- [39] Y. Kim, P.M. Thompson, L.A. Vese, HARDI data denoising using vectorial total variation and logarithmic barrier, *Inverse Probl. Imag.* 4 (2) (2010) 273–310.
- [40] K.T. Block, M. Uecker, J. Frahm, Suppression of MRI truncation artifacts using total variation constrained data extrapolation, *Int J Biomed Imag.* 2008 (2008).
- [41] C. Gnahn, M. Bock, P. Bachert, W. Semmler, N.G.R. Behl, A.M. Nagel, Iterative 3D projection reconstruction of ^{23}Na data with ^1H MRI constraints, *Magnet. Reson. Med.* 71 (5) (2014) 1720–1732.

- [42] F. Huang, Y. Chen, G.R. Duensing, J. Akao, A. Rubin, C. Saylor, Application of partial differential equation-based inpainting on sensitivity maps, *Magnet. Reson. Med.* 53 (2) (2005) 388–397.
- [43] J. Kornak, K. Young, B. Soher, A. Maudsley, Bayesian k-space-time reconstruction of MR spectroscopic imaging for enhanced resolution, *IEEE Trans. Med. Imag.* 29 (7) (2010) 1333–1350.
- [44] M.A. Griswold, P.M. Jakob, R.M. Heidemann, M. Nittka, V. Jellus, J. Wang, B. Kiefer, A. Haase, Generalized autocalibrating partially parallel acquisitions (GRAPPA), *Magnet. Reson. Med.* 47 (6) (2002) 1202–1210.
- [45] J. Sánchez-González, J. Tsao, U. Dydak, M. Desco, P. Boesiger, K. Paul Pruessmann, Minimum-norm reconstruction for sensitivity-encoded magnetic resonance spectroscopic imaging, *Magnet. Reson. Med.* 55 (2) (2006) 287–295.
- [46] R. Otazo, F.-H. Lin, G. Wiggins, R. Jordan, D. Sodickson, S. Posse, Superresolution parallel magnetic resonance imaging: application to functional and spectroscopic imaging, *NeuroImage* 47 (1) (2009) 220–230.
- [47] S. Posse, G. Tedeschi, R. Risinger, R. Ogg, D. Le Bihan, High speed 1H spectroscopic imaging in human brain by echo planar spatial-spectral encoding, *Magn Reson Med* 33 (1) (1995) 34–40.
- [48] T. Kirchner, A. Fillmer, J. Tsao, K.P. Pruessmann, A. Henning, Reduction of voxel bleeding in highly accelerated parallel ¹H MRSI by direct control of the spatial response function, *Magnet. Reson. Med.* 73 (2) (2015) 469–480.
- [49] S. Bhawe, R. Eslami, M. Jacob, Sparse spectral deconvolution algorithm for noncartesian MR spectroscopic imaging, *Magnet. Reson. Med.* 71 (2) (2014) 469–476.
- [50] Z.-P. Liang, Spatiotemporal imaging with partially separable functions, in: 4th IEEE International Symposium on Biomedical Imaging: From Nano to Macro, 2007, ISBI 2007, April 2007, pp. 988–991.
- [51] M. Lustig, J. Santos, D. Donoho, J. Pauly, k-t SPARSE: high frame rate dynamic MRI exploiting spatio-temporal sparsity, in: Proceedings of the 14th Annual Meeting of ISMRM, Seattle, Washington, USA, 2006, p. 2420.
- [52] J. Haldar, Z.-P. Liang, Spatiotemporal imaging with partially separable functions: a matrix recovery approach, in: 2010 IEEE International Symposium on Biomedical Imaging: From Nano to Macro, April 2010, pp. 716–719.
- [53] H. Pedersen, S. Kozerke, S. Ringgaard, K. Nehrke, W.Y. Kim, k-t PCA: temporally constrained k-t BLAST reconstruction using principal component analysis, *Magnet. Reson. Med.* 62 (2009) 706–716.
- [54] S. Lingala, Y. Hu, E. DiBella, M. Jacob, Accelerated dynamic MRI exploiting sparsity and low-rank structure: k-t SLR, *IEEE Trans. Med. Imag.* 30 (5) (2011) 1042–1054.
- [55] M. Usman, C. Prieto, T. Schaeffter, P.G. Batchelor, k-t group sparse: a method for accelerating dynamic MRI, *Magnet. Reson. Med.* 66 (2011) 1163–1176.
- [56] P.J. Shin, P.E.Z. Larson, M.A. Ohliger, M. Elad, J.M. Pauly, D.B. Vigneron, M. Lustig, Calibrationless parallel imaging reconstruction based on structured low-rank matrix completion, *Magnet. Reson. Med.* 72 (4) (2014) 959–970.
- [57] H. Nguyen, X. Peng, M. Do, Z.-P. Liang, Denoising MR spectroscopic imaging data with low-rank approximations, *IEEE Trans. Biomed. Eng.* 60 (1) (2013) 78–89.
- [58] I.S. Iohvidov, *Hankel and Toeplitz Matrices and Forms*, first ed., Birkhäuser, Boston, 1982.
- [59] J. Kasten, F. Lazeyras, D. Van De Ville, Data-driven MRSI spectral localization via low-rank component analysis, *IEEE Trans. Med. Imag.* 32 (10) (2013) 1853–1863.
- [60] K. Bredies, K. Kunisch, T. Pock, Total generalized variation, *SIAM J. Imag. Sci.* 3 (3) (2010) 492–526.
- [61] T. Chan, A. Marquina, P. Mulet, High-order total variation-based image restoration, *SIAM J. Sci. Comput.* 22 (2) (2000) 503–516.
- [62] S. Setzer, G. Steidl, Variational methods with higher order derivatives in image processing, in: M. Neamtu, L. Schumaker (Eds.), *Approximation XII*, Nashboro Press, Brentwood, 2008, pp. 360–386.
- [63] G. Maso, I. Fonseca, G. Leoni, M. Morini, A higher order model for image restoration: the one-dimensional case, *SIAM J. Math. Anal.* 40 (6) (2009) 2351–2391.
- [64] S. Lefkimiatis, J. Ward, M. Unser, Hessian Schatten-norm regularization for linear inverse problems, *IEEE Trans. Image Process.* 22 (5) (2013) 1873–1888.
- [65] T. Chan, S. Esedoglu, F. Park, A. Yip, Recent developments in total variation image restoration, in: N. Paragios, Y. Chen, O.D. Faugeras (Eds.), *Handbook of Mathematical Models in Computer Vision*, Springer-Verlag, 2005.
- [66] S. Lingala, M. Jacob, A blind compressive sensing frame work for accelerated dynamic MRI, in: 2012 9th IEEE International Symposium on Biomedical Imaging (ISBI), 2012, pp. 1060–1063.
- [67] K. Kreutz-Delgado, J.F. Murray, B.D. Rao, K. Engan, T.-W. Lee, T.J. Sejnowski, Dictionary learning algorithms for sparse representation, *Neural Comput.* 15 (2) (2003) 349–396.
- [68] M. Yaghoobi, T. Blumensath, M. Davies, Dictionary learning for sparse approximations with the majorization method, *IEEE Trans. Signal Process.* 57 (6) (2009) 2178–2191.
- [69] F. Lam, Z.-P. Liang, A subspace approach to high-resolution spectroscopic imaging, *Magnet. Reson. Med.* 71 (4) (2014) 1349–1357.
- [70] E. Adalsteinsson, P. Irarrazabal, S. Topp, C. Meyer, A. Macovski, D.M. Spielman, Volumetric spectroscopic imaging with spiral-based k-space trajectories, *Magnet. Reson. Med.* 39 (6) (1998) 889–898.
- [71] D.-H. Kim, R. Henry, D.M. Spielman, Fast multivoxel two-dimensional spectroscopic imaging at 3 T, *Magnet. Reson. Imag.* 25 (8) (2007) 1155–1161.
- [72] S. Sarkar, K. Heberlein, X. Hu, Truncation artifact reduction in spectroscopic imaging using a dual-density spiral k-space trajectory, *Magnet. Reson. Imag.* 20 (10) (2002) 743–757.
- [73] Z.-H. Wu, F. Lam, C. Ma, and Z.-P. Liang, Improved image reconstruction for subspace-based spectroscopic imaging using non-quadratic regularization, in: *Engineering in Medicine and Biology Society (EMBC), 2014 36th Annual International Conference of the IEEE*, 2014, pp. 2432–2435.
- [74] P. Cao, E.X. Wu, Accelerating phase-encoded proton MR spectroscopic imaging by compressed sensing, *J. Magn. Reson. Imag.* 41 (2) (2015) 487–495.
- [75] R. Otazo, D. Sodickson, A. Yoshimoto, S. Posse, Accelerated proton echo-planar spectroscopic imaging using parallel imaging and compressed sensing, in: *Proc. Intl. Soc. Mag. Reson. Med* 17, 2009, p. 331.
- [76] N. Askin, B. Atis, E. Ozturk-Isik, Accelerated phosphorus magnetic resonance spectroscopic imaging using compressed sensing, in: *Engineering in Medicine and Biology Society (EMBC), 2012 Annual International Conference of the IEEE*, 2012, pp. 1106–1109.
- [77] S. Hu, M. Lustig, A.P. Chen, J. Crane, A. Kerr, D.A. Kelley, R. Hurd, J. Kurhanewicz, S.J. Nelson, J.M. Pauly, D.B. Vigneron, Compressed sensing for resolution enhancement of hyperpolarized ¹³C flyback 3D-MRSI, *J. Magnet. Reson.* 192 (2) (2008) 258–264.
- [78] V. Morozov, *Methods for solving incorrectly posed problems*, in: Z. Nashed (Ed.), first ed., Springer-Verlag, 1984.
- [79] P.C. Hansen, *Rank-Deficient and Discrete Ill-Posed Problems: Numerical Aspects of Linear Inversion*, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 1998.
- [80] G.H. Golub, M.T. Heath, G. Wahba, Generalized cross-validation as a method for choosing a good ridge parameter, *Technometrics* 21 (1979) 215–223.
- [81] P.C. Hansen, Analysis of discrete ill-posed problems by means of the L-curve, *SIAM Rev.* 34 (4) (1992) 561–580.
- [82] A. Thompson, J. Brown, J. Kay, D. Titterton, A study of methods of choosing the smoothing parameter in image restoration by regularization, *IEEE Trans. Pattern Anal. Mach. Intell.* 13 (4) (1991) 326–339.
- [83] N. Galatsanos, A. Katsaggelos, Methods for choosing the regularization parameter and estimating the noise variance in image restoration and their relation, *IEEE Trans. Image Process.* 1 (3) (1992) 322–336.
- [84] R. Tibshirani, Regression shrinkage and selection via the lasso, *J. Roy. Stat. Soc. Ser. B* 58 (1994) 267–288.
- [85] H. Liao, F. Li, M.K. Ng, Selection of regularization parameter in total variation image restoration, *J. Opt. Soc. Am. A* 26 (11) (2009) 2311–2320.
- [86] J.P. Oliveira, J.M. Bioucas-Dias, M.A. Figueiredo, Adaptive total variation image deblurring: a majorization–minimization approach, *Signal Process.* 89 (9) (2009) 1683–1693.
- [87] S. Babacan, R. Molina, A. Katsaggelos, Parameter estimation in TV image restoration using variational distribution approximation, *IEEE Trans. Image Process.* 17 (3) (2008) 326–339.
- [88] C.M. Stein, Estimation of the mean of a multivariate normal distribution, *Ann. Stat.* 9 (6) (1981) 1135–1151.
- [89] T. Blu, F. Luisier, The SURE-LET approach to image denoising, *IEEE Trans. Image Process.* 16 (11) (2007) 2778–2786.
- [90] J. Pesquet, A. Benazza-Benyahia, C. Chau, A SURE approach for digital signal/image deconvolution problems, *IEEE Trans. Signal Process.* 57 (12) (2009) 4616–4632.
- [91] S. Ramani, Z. Liu, J. Rosen, J. Nielsen, J. Fessler, Regularization parameter selection for nonlinear iterative image restoration and MRI reconstruction using GCV and SURE-Based methods, *IEEE Trans. Image Process.* 21 (8) (2012) 3659–3672.
- [92] J. Rissanen, Modeling by shortest data description, *Automatica* 14 (5) (1978) 465–471.
- [93] H. Akaike, A new look at the statistical model identification, *IEEE Trans. Autom. Control* 19 (6) (1974) 716–723.
- [94] R.B. Cattell, The scree test for the number of factors, *Multivar. Behav. Res.* 1 (2) (1966) 245–276.
- [95] D.W. Ostens, Selection of optimal regression models via cross-validation, *J. Chemometr.* 2 (1) (1988) 39–48.
- [96] S. Wold, Cross-validated estimation of the number of components in factor and principal components models, *Technometrics* 20 (4) (1978) 397–405.
- [97] G. Diana, C. Tommasi, Cross-validation methods in principal component analysis: a comparison, *Stat. Methods Appl.* 11 (1) (2002) 71–82.
- [98] G. Schwarz, Estimating the dimension of a model, *Ann. Stat.* 6 (2) (1978) 461–464.
- [99] R.E. Kass, A.E. Raftery, Bayes Factors and Model Uncertainty, Department of Statistics, University of Washington, Tech. Rep. 254, March 1993.
- [100] M.E. Tipping, C.M. Bishop, Probabilistic principal component analysis, *J. Roy. Stat. Soc. Ser. B (Stat. Methodol.)* 61 (3) (1999) 611–622.
- [101] T. Minka, Automatic Choice of Dimensionality for PCA, M.I.T., MIT Media Laboratory, Vision and Modeling Group 20 Ames Street, Cambridge, MA 02139, Tech. Rep. 514, December 2000.
- [102] D.C. Hoyle, Automatic PCA dimension selection for high dimensional data and small sample sizes, *J. Mach. Learn. Res.* 9 (2008) 2733–2759.
- [103] T. Hastie, R. Tibshirani, J. Friedman, *The Elements of Statistical Learning*, Springer Series in Statistics, Springer New York Inc., New York, NY, USA, 2001.
- [104] D. Donoho, Compressed sensing, *IEEE Trans. Inform. Theory* 52 (4) (2006) 1289–1306.
- [105] M. Lustig, D. Donoho, J. Santos, J. Pauly, Compressed sensing MRI, *Signal Process. Mag. IEEE* 25 (2) (2008) 72–82.

- [106] A. Maudsley, C. Domenig, V. Govind, A. Darkazanli, C. Studholme, K. Arheart, C. Bloomer, Mapping of brain metabolite distributions by volumetric proton MR spectroscopic imaging (MRSI), *Magnet. Reson. Med.* 61 (3) (2009) 548–559.
- [107] W. Bogner, S. Gruber, S. Trattng, M. Chmelik, High-resolution mapping of human brain metabolites by free induction decay 1H MRSI at 7 T, *NMR Biomed.* 25 (6) (2012) 873–882.
- [108] D.-C. Woo, B.-S. Kim, S.-L. Jung, H.-J. Park, H.-S. Rhim, G.-H. Jahng, B.-Y. Choe, Development of a cone-shape phantom for multi-voxel MR spectroscopy, *J. Neurosci. Methods* 162 (2007) 101–107.
- [109] J.A. Kasten, T. Vetterli, F. Lazeyras, D. Van De Ville, 3D-printed shepp-logan phantom as a real-world benchmark for MRI, *Magnet. Reson. Med.* (2015) (Epub ahead of print).