

Separation of Metabolites and Macromolecules for Short-TE ¹H-MRSI Using Learned Component-Specific Representations

Journal:	<i>IEEE Transactions on Medical Imaging</i>
Manuscript ID	TMI-2020-1985.R1
Manuscript Type:	Full Paper
Date Submitted by the Author:	n/a
Complete List of Authors:	Li, Yahang; University of Illinois at Urbana-Champaign, Bioengineering Wang, Zepeng; University of Illinois at Urbana-Champaign, Bioengineering Sun, Ruoyu; University of Illinois at Urbana-Champaign, Industrial and Enterprise Systems Engineering Lam, Fan; University of Illinois at Urbana-Champaign, Beckman Institute
Keywords:	Magnetic resonance imaging (MRI) < Imaging modalities, Brain < Object of interest, Molecular and cellular imaging < Imaging modalities, Inverse methods < General methodology, Machine learning < General methodology, Image reconstruction - iterative methods < General methodology

Separation of Metabolites and Macromolecules for Short-TE ^1H -MRSI Using Learned Component-Specific Representations

Yahang Li, Zepeng Wang, Ruoyu Sun, Fan Lam

Abstract—Short-echo-time (TE) proton magnetic resonance spectroscopic imaging (MRSI) allows for simultaneously mapping a number of molecules in the brain, and has been recognized as an important tool for studying in vivo biochemistry in various neuroscience and disease applications. However, separation of the metabolite and macromolecule (MM) signals present in the short-TE data with significant spectral overlaps remains a major technical challenge. This work introduces a new approach to solve this problem by integrating imaging physics and representation learning. Specifically, a mixed unsupervised and supervised learning-based strategy was developed to learn the metabolite and MM-specific low-dimensional representations using deep autoencoders. A constrained reconstruction formulation is proposed to integrate the MRSI spatio-spectral encoding model and the learned representations as effective constraints for signal separation. An efficient algorithm was developed to solve the resulting optimization problem with provable convergence. Simulation and experimental results have been obtained to demonstrate the component-specific representation power of the learned models and the capability of the proposed method in separating metabolite and MM signals for practical short-TE ^1H -MRSI data.

Index Terms—proton (^1H) magnetic resonance spectroscopic imaging, short TE, signal separation, deep learning, deep autoencoder, low-dimensional models

I. INTRODUCTION

PROTON MRSI (^1H -MRSI) is a unique molecular imaging modality that can noninvasively map various endogenous metabolites in the brain. This molecular-level information has been demonstrated useful in different neuroscience and clinical applications, including brain tumors [1], [2], metabolic disorders [3], and neurodegenerative diseases [4], [5]. Short-echo-time (TE) ^1H -MRSI, in particular, offers several unique advantages compared to the more commonly used long-TE acquisitions, such as higher signal-to-noise ratio (SNR) due to less relaxation-induced signal loss and improved detection and quantification of molecules with short T_2 's and/or J -coupled

spins, e.g., myo-inositol (mI), glutamate (Glu) and glutamine (Gln) [6]–[10]. However, the applications of short-TE ^1H -MRSI have been limited by several technical challenges. One of these major challenges is the presence of macromolecule (MM) signals that overlap with the metabolite signals across the entire spectrum. This makes accurate and reproducible metabolite quantification difficult. It has been demonstrated that metabolite quantification can be substantially improved with better characterization and separation of MM signals from the short-TE data [11]–[14]. Moreover, the separated MM components may also provide additional biomarkers for various disease applications [13], [15]–[17].

A number of methods have been proposed to separate the metabolite and MM signals in short-TE MRSI data. One approach is to suppress the metabolite or MM signals during the data acquisition stage by exploiting their longitudinal relaxation (T_1) or diffusion property differences [18]. Examples include the most commonly used inversion recovery (IR) based excitation strategies, which are designed to null either the metabolite (with longer T_1 's) or MM (with shorter T_1 's) signals to measure the other [11], [19]–[22]. Methods that use two acquisitions to obtain both metabolites and MMs have also been proposed [13], [23]. Due to the variable ranges of T_1 values for different molecules in vivo, complete nulling of metabolites or MMs is impossible and additional processing are usually needed to further remove the residual spectral components. If both metabolites and MMs are desired, two acquisitions are needed which will inherently increase the imaging time.

An alternative signal processing based approach is to model the overall short-TE data using parametric models of metabolites and MMs individually. The separation is then achieved by estimating the model parameters for each component from the data, e.g., solving a constrained nonlinear least-square problem [11], [24]–[29]. Improved fitting strategies have been proposed to take advantage of the fast decaying nature of the MM signals for better separation. Specifically, one can fit and back-extrapolate the metabolites using a truncated FID with negligible MM contributions and estimate the MM component by subtracting the extrapolated metabolite fits from the original signal [27], [30], [31]. Iterative subtraction and refitting can be done for improved separation performance. However, these methods are sensitive to model mismatch and noise, and often lead to substantial voxel-to-voxel estimation variations for practical MRSI data. Nonetheless, inspired by

Correspondence to F. Lam, Ph.D., Department of Bioengineering and the Beckman Institute for Advanced Science and Technology, University of Illinois at Urbana-Champaign, Urbana, IL, 61801 USA e-mail: fan-lam1@illinois.edu.

Y. Li and Z. Wang are with the Department of Bioengineering, University of Illinois at Urbana-Champaign.

R. Sun is with the Department of Industrial and Enterprise Systems Engineering, University of Illinois at Urbana-Champaign.

This work is supported in part by NSF-CBET-1944249 and NSF-CCF-1755847.

these parametric models, we recognized that the metabolites and MMs have their distinct spectral patterns specified by just a few spectral parameters, e.g., concentrations, resonance frequencies, and lineshapes, thus should reside in their own nonlinear low-dimensional manifolds embedded in the original high-dimensional space [32]. We hypothesized that these manifolds could be learned from specially designed training data and used as effective constraints for metabolite and MM separation.

While learning nonlinear low-dimensional models from high-dimensional heterogeneous data has been a major challenge in machine learning, recent breakthroughs in deep learning have enabled excellent solutions to such problems [33]–[35]. Leveraging this progress, deep neural networks have been successfully adapted to process and quantify short-TE MRSI data with MM separation/removal capability [36]–[39]. The initial attempts have been focused on training an end-to-end network that learns the inverse function to directly map the noisy and artifact containing data to the desired spectral parameters [36], [39], spectra with MM components removed [37], or even a group of networks to extract individual metabolites [38]. These methods require the complicated networks to simultaneously capture the physical model, metabolite and MM spectral variations, and all other nuances related to noise, artifacts, and acquisition designs. Recently, an alternative approach has been proposed to use deep networks to learn a low-dimensional representation of general MR spectra and use this as a prior in a constrained reconstruction formalism instead of a direct inverse mapping. This approach not only simplifies the learning problem but also allows for more flexible integration with the physical forward model for different acquisition designs. And it has been successfully applied to ^{31}P -MRSI reconstruction [40].

Inspired by this approach, we proposed here a novel method to separate metabolite and MM signals for short-TE ^1H -MRSI by learning their distinct nonlinear low-dimensional models and incorporating the learned models into a constrained reconstruction formulation. Specifically, we proposed a new strategy that combines supervised and unsupervised learning to train two special deep autoencoders (DAEs) to learn efficient low-dimensional representations that are specific to metabolites and MMs, respectively. We devised a formulation to integrate the learned representations as effective constraints with a spatio-spectral encoding model for joint reconstruction as well as signal separation. An efficient algorithm was developed to solve the resulting optimization problem. We demonstrated the efficient and component-specific low-dimensional representations learned by our DAEs for metabolites and MMs, respectively. Numerical simulations and in vivo experiments were performed to illustrate the superior separation performance achieved by the proposed method over the standard parametric fitting method. Theoretical convergence analysis for the proposed algorithm was also provided. The following sections describe the proposed model learning strategy, reconstruction formulation, and numerical algorithm in details.

II. BACKGROUND

A. Metabolite and Macromolecule Signal Separation

In a short-TE ^1H -MRSI acquisition, the data will contain non-negligible contributions from both a metabolite component $\rho_{met}(\mathbf{r}, t)$ and a macromolecule component $\rho_{MM}(\mathbf{r}, t)$. The goal of separating these two signal components can be mathematically defined as estimating $\rho_{met}(\mathbf{r}, t)$ and $\rho_{MM}(\mathbf{r}, t)$ from their summation:

$$\rho(\mathbf{r}, t) = \rho_{met}(\mathbf{r}, t) + \rho_{MM}(\mathbf{r}, t), \quad (1)$$

which is an ill-posed problem. Solving this problem requires effective constraints. One most common approach is to impose parametric models on each component, i.e., $\rho_{met}(\mathbf{r}, t) = f_{met}(t; \alpha(\mathbf{r}))$ and $\rho_{MM}(\mathbf{r}, t) = f_{MM}(t; \beta(\mathbf{r}))$ where $\alpha(\mathbf{r})$ and $\beta(\mathbf{r})$ contain spectral parameters for the metabolites and MMs, respectively. While $f_{met}(t; \alpha(\mathbf{r}))$ usually incorporates resonance structures/metabolite basis generated by quantum mechanical simulations or phantom measurements, models for $f_{MM}(t; \beta(\mathbf{r}))$ are generally considered to be less molecule-specific (some MM peaks can be attributed to specific amino-acid residues in various proteins which may help to generate potentially stronger spectral priors [19]). As a result, different models including polynomials [27], [41], splines [24], [28], wavelets [26], and Gaussian lineshapes [13], [21], [22], [29] have been considered for f_{MM} . Gaussian lineshape based models with a priori determined chemical shift frequencies (from extensive in vivo and in vitro experiments) have shown a great balance between model complexity and fitting accuracy.

Specifically, a widely used model for both metabolites and MMs for an individual FID in short-TE ^1H -MRSI data can be written as follows [14], [37], [42]

$$\begin{aligned} \rho(t) = & \sum_{m=1}^M c_m e^{i(\phi_0 + \phi_m)} v_m(t) e^{-t/T_{2,m}^* + i2\pi\delta f_m t} \\ & + \sum_{l=1}^L b_l e^{i(\phi_0 + \psi_l)} e^{-t^2 \frac{\pi^2 W_l^2}{4 \ln(2)} + i2\pi\delta f_l t}, \end{aligned} \quad (2)$$

where the first summation represents the metabolite signals with c_m , $T_{2,m}^*$, and δf_m denoting the concentrations coefficients, physiology/experiment-dependent lineshapes and frequency shifts for individual molecules, respectively, and $\{v_m(t)\}$ corresponds to the metabolite basis. The second term captures the MM signals where W_l and δf_l denote the Gaussian linewidths and resonance frequencies for each MM group. The variables ϕ_0 , ϕ_m , and ψ_l are a global zeroth-order phase and molecule-dependent phases. While directly estimating all these parameters from a single FID or spectrum can be rather challenging and lead to large estimation variances, these models imply that the metabolite and MM signals may reside in their own nonlinear low-dimensional manifolds, which we believe can be learned and then used as effective constraints for metabolite and MM separation.

III. PROPOSED METHOD

A. Learning Component-Specific Low-Dimensional Models

Learning a single low-dimensional model for MR spectra by treating the entire spectrum as a point in a high-dimensional space has been investigated in [40], [43], [44].




Fig. 1. Illustration of the proposed model learning strategy with a mixture of supervised and unsupervised learning. Specifically, two DAEs are designed to capture the metabolite and MM-specific low-dimensional representations. For metabolite DAE (metabolite-specific NN), the unsupervised part enforces the network to extract a set of low-dimensional features that can approximate the metabolite signals accurately while the supervised part uses zeros as labels for the corresponding MM inputs. This enforces the network to learn to minimize its representation power of the MM signals, which will be useful for signal separation. A similar training strategy is applied to MM-specific DAE (macromolecule-specific NN) with the roles of metabolites and MMs exchanged. The mathematical formulations of this training strategy are provided in Eqs. (3) and (4).

While these learned models are powerful for spatio-spectral reconstruction or denoising, they are not suited to address the signal separation problem, which requires component-specific constraints. The metabolite and MM signals have their own unique signal characteristics (e.g., distinct spectral features and parameter distributions); thus, separate models can be learned to capture the low-dimensional manifolds where they reside on or close to. However, a straightforward application of the previously described DAE [40] trained using metabolites and MMs individually may not be effective since they are not optimized to differentiate the two spectrally overlapping signal components. Therefore, we proposed here a mixed supervised and unsupervised learning strategy to address this issue. Specifically, we seek to train a DAE that can extract an accurate and efficient low-dimensional representation of the metabolite (or MM) signals with simultaneously minimal representation capability of the other component. Figure 1 provides a graphical illustration of this special model learning strategy. Mathematically, the learning problems are formulated as follows, for metabolites:

$$\begin{aligned} \{\hat{\theta}_{met}\} = \arg \min_{\theta_{met}} \frac{1}{N} \left[\sum_{n=1}^N \epsilon(\mathbf{x}_{met}^n - \mathcal{N}_{met}(\mathbf{x}_{met}^n; \theta_{met})) \right. \\ \left. + \|\mathcal{N}_{met}(\mathbf{x}_{mm}^n; \theta_{met})\|_2^2 \right], \end{aligned} \quad (3)$$

and for macromolecules:

$$\begin{aligned} \{\hat{\theta}_{mm}\} = \arg \min_{\theta_{mm}} \frac{1}{N} \left[\sum_{n=1}^N \epsilon(\mathbf{x}_{mm}^n - \mathcal{N}_{mm}(\mathbf{x}_{mm}^n; \theta_{mm})) \right. \\ \left. + \|\mathcal{N}_{mm}(\mathbf{x}_{met}^n; \theta_{mm})\|_2^2 \right] \end{aligned} \quad (4)$$

$\{\mathbf{x}_{met}^n\}_{n=1}^N$ and $\{\mathbf{x}_{mm}^n\}_{n=1}^N$ are training sets (FIDs) for the metabolites and MMs, respectively, with each $\mathbf{x}^n \in \mathcal{R}^{2T}$ being a sample with real and imaginary parts concatenated. N is the number of training samples and T is the length of the FID. $\mathcal{N}_{met}(\cdot, \cdot)$ and $\mathcal{N}_{mm}(\cdot, \cdot)$ denote the metabolite and MM component-specific neural networks (NNs) parameterized by θ_{met} and θ_{mm} , respectively. Here in the first term of Eq. (3), ϵ represents the loss for training that measures the error between the original metabolite input and the network approximation. This term enforces the NN to learn an accurate low-dimensional representation of the metabolite signals. The second term of Eq. (3), appearing as a regularization, is designed to minimize the output of the metabolite NN that corresponds to the MM data ($\mathbf{x}_{mm}^n$). This can also be viewed as using zeros as the labels for the MM input. As a result, the metabolite NN is trained to learn a representation specific to metabolite signal features with minimal representation power for MMs. Likewise, the two terms of Eq. (4) serve a similar purpose (to capture MM-specific low-dimensional features while minimizing metabolite representation). We hypothesized that DAEs trained separately in this fashion would have not only the ability to extract accurate nonlinear low-dimensional representations of metabolites and MMs individually but also the desired property that inaccurately models the other component. The specific network has an embedded "bottleneck" encoding-decoding structure that encodes the high-dimensional data into a set of L -dimensional features that can recover the original data, where L is referred to as the model order below. More details on the network are provided in the supplementary materials (Fig. S4).

B. Signal Separation Using the Learned Models

With the trained DAEs $\mathcal{N}_{met}(\cdot, \cdot)$ and $\mathcal{N}_{mm}(\cdot, \cdot)$ capturing learned metabolite and MM-specific models, one remaining challenge is to effectively utilize the learned models for signal separation from practical $^1\text{H-MRSI}$ data. To this end, we proposed a regularized reconstruction formulation that integrates the forward spatio-spectral encoding model with B_0 inhomogeneity correction capability and the two learned models for metabolite and MM separated reconstruction. Specifically, we formulated the separation problem as:

$$\begin{aligned} \min_{\mathbf{X}_{met} \in \mathcal{X}_1, \mathbf{X}_{mm} \in \mathcal{X}_2} & \|\mathbf{d} - \Omega\{\mathbf{FB} \odot (\mathbf{X}_{met} + \mathbf{X}_{mm})\}\|_2^2 \\ & + \lambda_1 \sum_{n=1}^N \|\mathcal{N}_{met}(\mathbf{X}_{met}^n) - \mathbf{X}_{met}^n\|_2^2 \\ & + \lambda_2 \sum_{n=1}^N \|\mathcal{N}_{mm}(\mathbf{X}_{mm}^n) - \mathbf{X}_{mm}^n\|_2^2 \\ & + \lambda_3 \|\mathbf{D}_w(\mathbf{X}_{met} + \mathbf{X}_{mm})\|_F^2, \end{aligned} \quad (5)$$

where $\mathbf{X}_{met} \in \mathcal{C}^{N \times T}$ and $\mathbf{X}_{mm} \in \mathcal{C}^{N \times T}$ are matrix representations of the spatiotemporal functions of interest for the metabolite and MM components, with each row being a T -point FID and N the number of voxels. The feasible sets $\mathcal{X}_1, \mathcal{X}_2$ are balls with large enough radius that contain ground-truth representations¹. $\mathbf{B}$ models the linear phases induced by B_0 inhomogeneity, $\odot$ represents a point-wise multiplication, $\mathbf{F}$ denotes the Fourier transform, Ω is a (k, t) -space sampling operator (allows for flexible sampling designs), and $\mathbf{d}$ is a vector containing the noisy measured data. The first term enforces the imaging model and data consistency. The next two terms impose the priors that FID signals of metabolites and MMs belong to their own low-dimensional manifolds captured by the learned DAEs. The last term is a spatial smoothness constraint with $\mathbf{D}_w$ being a weighted finite-difference operator [45], and $\|\cdot\|_F$ denoting the Frobenius norm. Eq. (5) results in a high-dimensional optimization problem, which is challenging to solve due to the presence of both nonlinear functions related to the DAEs and quadratic functions of $\mathbf{X}_{met}$ and $\mathbf{X}_{mm}$.

C. Optimization Algorithm

We developed an efficient algorithm to address the computational challenges associated with the problem in Eq. (5). Specifically, we introduced an auxiliary variable

$$\mathbf{S} = \mathbf{B} \odot (\mathbf{X}_{met} + \mathbf{X}_{mm}) \quad (6)$$

and reformulated the problem as:

$$\begin{aligned} & \{\hat{\mathbf{X}}_{met}, \hat{\mathbf{X}}_{mm}, \hat{\mathbf{S}}\} \\ & = \arg \min_{\mathbf{X}_{met}, \mathbf{X}_{mm}, \mathbf{S}} \|\mathbf{d} - \Omega\{\mathbf{FS}\}\|_2^2 \\ & + \lambda_1 \sum_{n=1}^N \|\mathcal{N}_{met}(\mathbf{X}_{met}^n) - \mathbf{X}_{met}^n\|_2^2 \\ & + \lambda_2 \sum_{n=1}^N \|\mathcal{N}_{mm}(\mathbf{X}_{mm}^n) - \mathbf{X}_{mm}^n\|_2^2 \\ & + \lambda_3 \|\mathbf{D}_w \bar{\mathbf{B}} \odot \mathbf{S}\|_F^2 \\ & \text{s.t. } \mathbf{B} \odot (\mathbf{X}_{met} + \mathbf{X}_{mm}) = \mathbf{S}, \\ & \mathbf{X}_{met} \in \mathcal{X}_1, \mathbf{X}_{mm} \in \mathcal{X}_2, \end{aligned} \quad (7)$$

where $\bar{\mathbf{B}}$ denotes element-wise conjugate of $\mathbf{B}$. Then, the alternating direction method of multipliers (ADMM) is adapted to solve this equivalent problem [46], in which it is decomposed into simpler linear least-squares problems and nonlinear problems that can be solved in a parallel fashion. More specifically, the following subproblems were solved iteratively:

(I) Update $\mathbf{X}_{met}$ with fixed $\mathbf{X}_{mm}^{(i)}$ and $\mathbf{S}^{(i)}$ as follows (i is the iteration index)

$$\begin{aligned} \mathbf{X}_{met}^{(i+1)} = \arg \min_{\mathbf{X}_{met} \in \mathcal{X}_1} & \lambda_1 \sum_{n=1}^N \|\mathcal{N}_{met}(\mathbf{X}_{met}^n) - \mathbf{X}_{met}^n\|_2^2 \\ & + \frac{\rho}{2} \left\| \mathbf{B} \odot (\mathbf{X}_{met} + \mathbf{X}_{mm}^{(i)}) - \mathbf{S}^{(i)} + \frac{\mathbf{Y}^{(i)}}{\rho} \right\|_F^2 \end{aligned} \quad (8)$$

where $\mathbf{Y}^{(i)}$ is the Lagrangian multiplier and ρ is the penalty parameter.

(II) Update $\mathbf{X}_{mm}$ with fixed $\mathbf{X}_{met}^{(i+1)}$ and $\mathbf{S}^{(i)}$ as

$$\begin{aligned} \mathbf{X}_{mm}^{(i+1)} = \arg \min_{\mathbf{X}_{mm} \in \mathcal{X}_2} & \lambda_2 \sum_{n=1}^N \|\mathcal{N}_{mm}(\mathbf{X}_{mm}^n) - \mathbf{X}_{mm}^n\|_2^2 \\ & + \frac{\rho}{2} \left\| \mathbf{B} \odot (\mathbf{X}_{met}^{(i+1)} + \mathbf{X}_{mm}) - \mathbf{S}^{(i)} + \frac{\mathbf{Y}^{(i)}}{\rho} \right\|_F^2 \end{aligned} \quad (9)$$

(III) Update $\mathbf{S}$ with $\mathbf{X}^{(i+1)} = \mathbf{X}_{met}^{(i+1)} + \mathbf{X}_{mm}^{(i+1)}$ by solving

$$\begin{aligned} \mathbf{S}^{(i+1)} = \arg \min_{\mathbf{S}} & \|\mathbf{d} - \Omega\{\mathbf{FS}\}\|_2^2 + \lambda_3 \|\mathbf{D}_w \bar{\mathbf{B}} \odot \mathbf{S}\|_F^2 \\ & + \frac{\rho}{2} \left\| \mathbf{B} \odot \mathbf{X}^{(i+1)} - \mathbf{S} + \frac{\mathbf{Y}^{(i)}}{\rho} \right\|_F^2 \end{aligned} \quad (10)$$

(IV) Update $\mathbf{Y}$ as

$$\mathbf{Y}^{(i+1)} = \mathbf{Y}^{(i)} + \rho(\mathbf{B} \odot \mathbf{X}^{(i+1)} - \mathbf{S}^{(i+1)}) \quad (11)$$

Subproblems (I) and (II) contain both terms associated with the nonlinear networks, $\mathcal{N}_{met}(\cdot, \cdot)$ and $\mathcal{N}_{mm}(\cdot, \cdot)$, and can be solved using a generic nonlinear optimization solver. Subproblem (III) is a typical linear least-squares problem with a quadratic regularization. Note that although directly minimizing Eq. (8) and Eq. (9) are very high-dimensional problems for which computing the gradient is very demanding, it can be solved in a voxel-by-voxel fashion since the Frobenius norm term is separable for all the voxels (i.e., individual rows in $\mathbf{X}$).

¹For the empirical computation, we often set $\mathcal{X}_1, \mathcal{X}_2$ to be $\mathcal{C}^{N \times T}$; but due to technical consideration, in the convergence results, we set them to be bounded sets.

Based on the autoencoder design, the gradients for individual voxels can be efficiently calculated through backpropagation. More specifically, denote

$$\begin{aligned} f_n(\mathbf{X}_{met}^n, \mathbf{X}_{mm}^n) = & \lambda_1 \|\mathcal{N}_{met}(\mathbf{X}_{met}^n) - \mathbf{X}_{met}^n\|_2^2 \\ & + \lambda_2 \|\mathcal{N}_{mm}(\mathbf{X}_{mm}^n) - \mathbf{X}_{mm}^n\|_2^2 \\ & + \frac{\rho}{2} \left\| \left[\mathbf{B} \odot \mathbf{X} - \mathbf{S}^{(i)} + \frac{\mathbf{Y}^{(i)}}{\rho} \right]_n \right\|_2^2 \end{aligned} \quad (12)$$

as the cost function for the n th voxel, the then gradients for the metabolite component can be written as:

$$\begin{aligned} \nabla f_n(\mathbf{X}_{met}^n) = & 2\lambda_1 (\mathbf{J}_{\mathcal{N}_{met}} - \mathbf{I})^T (\mathcal{N}_{met}(\mathbf{X}_{met}^n) - \mathbf{X}_{met}^n) \\ & + \rho \mathbf{B}_{(n)}^H \left(\mathbf{B}_{(n)} \mathbf{X}_n - \mathbf{S}_n^{(i)} + \frac{\mathbf{Y}_n^{(i)}}{\rho} \right) \end{aligned} \quad (13)$$

where $\mathbf{J}_{\mathcal{N}_{met}} \in \mathcal{R}^{T \times T}$ is the Jacobian of the metabolite network, $\mathbf{I}$ is a $T \times T$ identity matrix, and $\mathbf{B}_{(n)}$ represents a diagonal matrix formed by the n th row of $\mathbf{B}$. The gradients for MM component can be derived similarly (omitted due to space constraint). And the Jacobians $\mathbf{J}_{\mathcal{N}_{met}}$ and $\mathbf{J}_{\mathcal{N}_{mm}}$ can be calculated through backpropagation described in [40].

Subproblem (III) is equivalent to solving a system of linear equations with a spatial regularization on the overall spatiotemporal function (due to the way the auxiliary variable is introduced). The iteration is terminated until a specified iteration number is reached (e.g., 20) or the relative change between $\mathbf{X}_{met}^{(i)}$ and $\mathbf{X}_{met}^{(i+1)}$ is below a threshold (e.g., 10^{-4}).

D. Training Data Generation

One common issue of training deep neural networks is the requirement of a large number of high-quality training data. Strategies that combine spectral fitting models, quantum mechanical (QM) simulations and experimental data have been described in several literature to address this issue for spectral model learning [37], [40], [47]. Utilizing a similar strategy here, we generated metabolite and MM training data separately using the model in Eq. (2). For metabolites, the basis $\mathbf{v}_m(t)$ were generated from QM simulations using the NMRScopeB software [48], for both the FID (pulse and acquire) and semi-LASER excitation schemes with different TEs [7]. These molecule specific resonance structures can be assumed to be invariant with respect to different subjects. Meanwhile, the empirical distributions of the spectral parameters, i.e., c_m , $T_{2,m}^*$, and δf_m were estimated from literature values as well as fitting experimental high-SNR, low-resolution MRSI data from healthy volunteers [37], [43], [47]. The empirical distributions were fitted to parametric Gaussian distributions to allow for generating more randomly distributed parameters. The global zeroth-order phase was generated from a Gaussian distribution with mean zero and standard deviation of 25 degrees, and Gaussian distributed molecule dependent phases were also introduced to simulate more realistic signal variations (with mean zero and standard deviation of 10 degrees). Finally, the metabolite basis and parameters randomly sampled from these distributions were combined using the model in Eq. (2) to generate 300,000 ^1H MR spectra. Metabolites commonly

observed and quantified in ^1H -MRSI experiments are considered, i.e., N-acetylaspartate (NAA), creatine (Cr), choline (Cho), glutamate (Glu), glutamine (Gln), myo-inositol (mI), gamma-Aminobutyric acid (GABA), glutathione (GSH) and lactate (Lac). For MMs, another 300,000 training samples were generated using a similar procedure and the model in Eq. (2). 13 commonly reported MM resonances with mean δf_l 's equal to 0.9, 1.21, 1.38, 1.63, 2.01, 2.09, 2.25, 2.61, 2.96, 3.11, 3.67, 3.8, and 3.96 ppm were included. The parameters b_l (concentration coefficients) and W_l (linewidths) for different peak groups were assumed to follow Gaussian distributions. The mean values were acquired from [37] with standard deviations specified as 20% of the means to introduce relative peak variations. A global scaling factor was introduced to the MM coefficients to reflect experimentally observed metabolite-to-MM signal ratios. Finally, all data were normalized to the range of -1 to 1 for training.

E. Other Implementation Details

Among the 300,000 training data, 200,000 were used for training and 100,000 for testing. All the model parameters were upper and lower bounded based on our own and other published ^1H -MRSI data [37], [38]. Specifically, the $T_{2,m}^*$ values were lower bounded by 5 ms and upper bounded by 200 ms, and the c_m values were bounded between 0 and 2 with the mean NAA concentration being 1. The linewidth W_l values were bounded within the range of 5 to 70 Hz, and b_l 's were lower bounded by 0. The parameters were first generated, and the values outside these ranges were excluded. The spectral bandwidth (BW) was fixed at 2000 Hz. Similar to the DAE used in [40], the metabolite and MM networks have a fully-connected structure of $2T - 1000 - 250 - 100 - L - 100 - 250 - 1000 - 2T$. Both Tanh and ReLu units can be used in the nonlinear hidden layers except the middle linear layer (with similar performances). The results shown below were from ReLu. The learned network models were first evaluated with a range of L 's. The exact L for phantom and in vivo data processing was chosen, such that the NNs for metabolites and MMs have similar approximation errors (around 5%). Note that the training only needs to be performed once for a fixed excitation scheme (with specific choices of RF pulses, TE, and field strength). All the networks were implemented in PyTorch and trained using an NVIDIA RTX Titan graphics processing unit on Windows 10 using the Adam optimizer [49] with a batch size of 500, an initial learning rate of 0.001, and 300 epochs while the other parameters remained as default. The Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm was used to solve the optimization problem for individual voxels in Eqs. (8) and (9), and the linear conjugate gradient was used to solve Eq. (10) [50].

IV. SIMULATION AND EXPERIMENTAL SETTINGS

A. Numerical Simulations

The component-specific representation power of our learned models was first evaluated. Specifically, we validated the approximation accuracy of the trained DAE-based metabolite and MM-specific low-dimensional models with comparison

to linear subspace models (estimated from the same training data) [43], at different model orders. Testing metabolite and MM data were generated and passed through the trained metabolite and MM-specific networks, respectively. The errors of the same test data projected onto the metabolite and MM subspaces were also calculated. More specifically, two Casorati matrices were first constructed by stacking all the training data for metabolites or MMs, respectively. Then the component-specific subspace was obtained by SVD with a rank truncation (L , the model order). Finally, the testing data were projected onto the two subspaces separately to evaluate approximation accuracy [40], [43]. The approximation performance was evaluated quantitatively using a relative ℓ_2 error defined as:

$$\text{err} = \frac{\|\mathbf{X}_{\text{true}} - \hat{\mathbf{X}}\|_F}{\|\mathbf{X}_{\text{true}}\|_F} \quad (14)$$

where $\mathbf{X}_{\text{true}}$ denotes the original data (each column being an FID), and $\hat{\mathbf{X}}$ represents the model approximation or reconstructed data (see below).

A numerical phantom was constructed to evaluate the signal separation performance using the learned DAE-based nonlinear models (details of the phantom generation process can be found in the supplementary materials). In short, brain tissue fraction maps for gray matter (GM), white matter (WM), and cerebrospinal fluid (CSF) were first obtained from an in vivo anatomical T_1 -weighted image. Then regional spectral parameters for different ^1H metabolites and MM components described in the Training Data Generation section were assigned based on literature values [14], [37], [51]. The constant parameters in each region were subsequently combined using tissue fraction maps as weightings to simulate continuously varying parameters across the brain. Finally, the parameters at different voxels along with metabolite basis $\{v_m\}$ were fed into Eq. (2) to synthesize spatially localized FIDs. To simulate a more realistic scenario, voxel-dependent random frequency shifts (mean zero and standard deviation of 5 Hz) for different molecules, as well as B_0 inhomogeneity (mean zero and standard deviation of 10 Hz), were also introduced. A lesion-like feature with significantly altered metabolite ratios (e.g., a factor of three higher Cho and lower concentrations for other metabolites) and a higher MM level was included. Noisy data were generated by adding complex white Gaussian noise with different SNRs to the simulated (x, t) -space data. The SNR is defined with respect to the maximum NAA peak amplitude within the FOV. After the proposed reconstruction, the separated metabolite and MM components were fitted individually using a metabolite-only and a MM-only parametric model from Eq. (2) in a voxel-by-voxel fashion to produce molecular maps. In comparison, a direct parametric fitting was also performed to the original data without the proposed separation. An FID truncation was performed to fit the metabolites first which were then back-extrapolated for subtraction to fit the MMs. A metabolite refitting was done after subtracting the fitted MMs from the original data. All fittings were done using in-house implementations which have been validated against the time-domain fitting using the jMRUI package [27], [48] (The customized implementations provided more flexibility for further optimizations).




Fig. 2. Representation capability of the learned nonlinear models: a) Approximation errors (relative ℓ_2) of the trained metabolite DAE (orange curve) compared to a linear subspace model (blue curve) for the metabolite data at different L 's; b) Approximation errors of the MM DAE (orange curve) with comparison to a linear subspace model (blue curve); c) and e) Representative metabolite and MM spectra (black), and the approximations of each signal by the metabolite DAE (orange) and metabolite linear subspace (blue) both with $L = 24$; d) and f) Representative MM and metabolite spectra (black), and the approximations by the MM DAE (orange) and subspace (blue) with $L = 8$. It is evident that the learned DAEs have more accurate and component-specific representation than the linear subspaces.

B. In Vivo Experiments

We have evaluated the performance of the proposed method using practical in vivo data acquired from five healthy volunteers with approval from the local Institutional Review Board. Experimental brain MRSI data were acquired on a 3T Prisma scanner equipped a 20-channel head coil using both an FID-MRSI (pulse and acquire) sequence and a semi-LASER MRSI short-TE sequence (sLASER). We chose these two sequences because they are among the most commonly used short-TE data acquisition schemes and require rather different metabolite basis sets which serves to demonstrate that the proposed method can be flexibly adapted to work with any sequences. The parameters for the FID-MRSI sequence were as follows: TR/TE = 800/4 ms, field-of-view (FOV) = $230 \times 230 \text{ mm}^2$, slice thickness = 10 mm, matrix size = 36×36 , spectral bandwidth (BW) = 2000 Hz and 512 FID samples. The total acquisition time was about 13.5 minutes with elliptical sampling. The parameters for the sLASER sequence were: TR/TE = 1600/40 ms, FOV = $180 \times 190 \text{ mm}^2$, slice thickness = 15 mm, matrix size = 24×24 , 2000 Hz BW and 1024 FID samples. The total acquisition time was about 16 minutes. A 60 Hz weak water suppression and carefully placed outer volume suppression bands were used for all the scans. Before reconstruction, the nuisance water and lipid signals were first removed using the method in [52] followed by coil combination of the water/lipid-removed data.




Fig. 3. Simulation results: a) Spatial variations of the overall metabolite and MM signals in the phantom (ℓ_2 integral along the FID dimension); b) A sampled voxel spectrum and its noisy counterpart (SNR=30); c)-e) Separation results from the proposed (orange curves) method and the direct parametric fitting (blue curves, without the proposed separation) for three different voxels (in GM, WM, and Lesion). The voxel locations are indicated by different shapes in (a). Similar overall spectra (c) were produced by both methods. But the proposed method yielded more accurate separated metabolite (d) and MM (e) components. The black arrows identify some spectral features better recovered by the proposed method.




Fig. 4. Simulation results: molecular maps of NAA, Cr, Cho, and MM from the ground truth (Gold Standard, column 1), the direct parametric fitting method (column 2) and the proposed method (column 3) are compared. For the proposed method, the maps were obtained by fitting the separated metabolite and MM components individually. The first MM peak group (located at ~ 0.9 ppm) is shown [37]. Note that the MM maps were normalized separately, thus having a different scaling compared to the metabolite maps. Relative ℓ_2 errors for the separated metabolite and MM signals are also calculated (shown in the images). The improved signal separation offered by the proposed method led to significantly improved molecular quantification.

V. RESULTS

A. Simulation Results

Figure 2 compares the representation powers of our learned component-specific nonlinear low-dimensional and the linear subspace models. As shown, the learned metabolite (Fig. 2a) and MM (Fig. 2b) DAE-based models achieved higher accu-

racy with lower relative ℓ_2 errors than the subspace models for their respective components. With the same model order (L), the MM DAE has higher accuracy than the metabolite DAE. The approximations for two representative testing metabolite and MM spectra are also shown in Fig. 2c-f to further demonstrate the accuracy and specificity of the learned models. With a fixed model order $L = 24$ for metabolite and $L = 8$ for MM, the metabolite DAE can accurately capture the metabolite spectral features (Fig. 2c), exhibiting a higher accuracy than the subspace model with the same dimension. Similar results can be observed for the MM test spectrum (Fig. 2d). More importantly, the metabolite DAE offers a poor approximation of the MM spectrum as we designed it to, while the linear metabolite subspace can still capture a decent amount of MM spectral energy (Fig. 2e), implying a weaker capability for signal separation. This component-specific representation can also be observed for the MM DAE (Fig. 2f), which does not capture the metabolite spectral features, while the MM subspace can again capture a large portion of metabolite signal energy. These validate the desirable component-specific representation capability of the learned models and imply their unique potential for improved metabolite and MM separation.

A set of metabolite and MM signal separation results from the numerical phantom (SNR=30) are shown in Figs. 3 and 4. Here the resulting spectra were shown in magnitude for visualization purpose (the real parts of reconstructed spectra can be found in the supplementary materials Fig. S1). The model order (L) was chosen as 24 for the metabolite DAE and 16 for the MM DAE with similar approximation errors ($\sim 5\%$ error), which achieved a good balance between model complexity and approximation accuracy. The regularization parameters λ_1 and λ_2 were chosen based on a single voxel separation performance, and λ_3 was chosen using the discrepancy principle and then fine-tuned by minimizing the ℓ_2 errors of the final spatiotemporal reconstructions. A time-domain direct parametric model-based fitting with back-extrapolation was also performed as described above, and the results were compared. The separated metabolite and MM spatiotemporal distributions from the proposed method were subject to parametric fitting




Fig. 5. Experimental results from the in vivo FID-MRSI data: a) and b) Maps of NAA, Cr, Cho, Glu and MM estimated from the separated signals produced by the proposed method (a) and from the direct parametric fitting method (b). The molecular maps were overlaid on an anatomical image for the matched slice; c) Spatially-resolved spectra from the voxel marked by the blue symbol, with the first and second rows showing the results from the proposed method (orange curves) and parametric fitting (blue curves), respectively. The original spectra are shown in black. The overall reconstruction, as well as the separated metabolite and MM spectra, are compared. The white arrows indicate some artifact-like features generated by the direct parametric fitting which are not present in the maps produced by the proposed method.




Fig. 6. Results from the same data in Fig. 5 but with the first 36 FID points truncated (~ 18 ms) and zero-padded to the original length. Two representative spatially-resolved spectra from the locations marked by the corresponding symbols are shown, including the original spectra (black), the overall reconstruction, and the separated metabolite and MM spectra (orange). Significantly reduced MM signals were observed, indicating that the proposed method is not overfitting.

(using the metabolite-only and MM-only parametric models, respectively). As can be seen, both the parametric fitting and the proposed method achieved similar estimates of the overall spectra (Fig. 3c), but the proposed method produced significantly more accurate separated metabolite and MM components (Figs. 3d and e). The metabolite maps obtained by direct parametric fitting of the overall data (Parametric Fitting) and the proposed method (separate fitting) are compared in Fig. 4. The results demonstrate the benefits of the proposed signal separation. More specifically, the metabolite maps from fitting the separated signals exhibit significantly less spatially dependent estimation variances and higher accuracy than those produced by a direct parametric fitting of the combined signals. An additional set of results from less noisy data (SNR=60) are shown in the supplementary materials (Fig. S3).

B. In Vivo Results

A set of spatially-resolved spectral reconstruction obtained by the proposed method from the in vivo FID-MRSI data are shown in Fig. 5. As can be seen, the proposed method was able to separate the metabolite and MM spectral components with a similar overall spectrum to the direct parametric fitting

method. Furthermore, the metabolite and MM maps from the proposed method exhibited a higher quality with less spatial estimation variances. More specifically, the molecule maps produced by the direct parametric fitting had apparent artifacts, e.g., locally dark/bright areas and sudden discontinuities (indicated by the white arrows in Fig. 5b). These artifacts were effectively reduced in the maps from the proposed method. The relative peak intensities may not appear the same as those from standard IR-based MM measurements due to different T_1 weightings (effects of no IR and the shorter TR used), e.g., a strong 0.9 ppm MM peak even before separation. Hence, results from another dataset acquired with a longer TR (1500 ms) are included in the supplementary materials (Fig. S2) to further illustrate the TR effects. We have also performed a reconstruction of the same FID-MRSI data with the first 36 time points truncated (i.e., 18 ms) to evaluate the robustness of the proposed method. As shown in Fig. 6, significantly reduced MM signals were obtained, indicating that the learned model is not overfitting.

Figure 7 shows the spatially-resolved spectral reconstruction from the sLASER data to further demonstrate the utility of the proposed method. Signals can only be observed from the central region of the brain due to the volume selective excitation. The proposed method again produced visually better separation, which can be observed in the metabolite and MM maps (e.g., better gray/white matter contrast and fewer artifacts) as well as the selective voxel spectra. The proposed method effectively reduced the over and under-estimation of some metabolite and MM spectral components in the parametric fitting method (Figs. 7b and d). Additional metabolite maps can be found in the supplementary materials (Fig. S5). Our approach should work for any excitation as long as the corresponding metabolite basis can be obtained for training data generation.

C. Convergence Analysis

We have also performed convergence analysis of the proposed algorithm. Figure 8 shows the relative changes between iterates ($\|\mathbf{X}^{(i+1)} - \mathbf{X}^{(i)}\|/\|\mathbf{X}^{(i)}\|$) and relative ℓ_2 errors for the metabolite and MM estimates w.r.t. the iteration number.




Fig. 7. Experimental results from the in vivo sLASER data: a) Maps of NAA, Cr and MM from the proposed method (overlaid on anatomical images); b) Two representative spatially-resolved spectra (voxel locations marked by the corresponding symbols) with the original spectra (black, column 1), the overall reconstruction (column 1), and the separated metabolite (column 2) and MM (column 3) spectra; c) and d) The corresponding results from the parametric fitting method with the same arrangement. The spectra from the two methods (b and d) are from the same voxels. The arrows indicate some over and underestimation of metabolite or MM components from the parametric fitting that were mitigated by the proposed method.

Empirical convergence can be observed. Furthermore, our problem formulation and the ADMM-based algorithm allow us to theoretically characterize its convergence.

Theorem 1 *There exists a constant ρ_0 such that if $\rho \geq \rho_0$, every limit point of the sequence $(\mathbf{X}^{(i)}, \mathbf{S}^{(i)})$ generated by the algorithm described in (8), (9), (10), and (11) is a stationary solution of the optimization problem (7) (i.e. a solution that satisfies the KKT condition).*

This theorem states that for a proper choice of penalty parameter ρ , the sequence generated by our algorithm is guaranteed to converge to stationary points. For non-convex problems, convergence to the global minimum is often very difficult. Thus, we follow the common practice to prove a result of convergence to stationary points. We remark that even the convergence to a stationary solution is a non-trivial property




Fig. 8. Convergence analysis of the algorithm: (a) Relative changes (in terms of ℓ_2 errors) between different iterates; (b) Relative ℓ_2 errors for the metabolite (orange) and MM (blue) components w.r.t. different iterations. As can be seen, the result changes minimally after 10 iterations.

because a general convergence result of ADMM for non-convex problems is still an open question. This theoretical characterization is enabled by the unique structure of our problem formulation: we show that it is a special case of the non-convex sharing problem [53], for which the convergence results have been established.

Proof sketch: If we denote $\mathbf{X}_1 = \text{vec}(\mathbf{X}_{met})$, $\mathbf{X}_2 = \text{vec}(\mathbf{X}_{mm})$, $\mathbf{X}_0 = \text{vec}(\mathbf{S})$, then $\lambda_1 \sum_{n=1}^N \|\mathcal{N}_{met}(\mathbf{X}_{met}^n) - \mathbf{X}_{met}^n\|_2^2$ is a function of $\mathbf{X}_1$, which we denote as $g_1(\mathbf{X}_1)$. Similarly, $\lambda_2 \sum_{n=1}^N \|\mathcal{N}_{mm}(\mathbf{X}_{mm}^n) - \mathbf{X}_{mm}^n\|_2^2$ is a function of $\mathbf{X}_2$, which we denote as $g_2(\mathbf{X}_2)$. The remaining terms, $\|\mathbf{d} - \Omega\{\mathbf{F}\mathbf{X}_0\}\|_2^2 + \lambda_3 \|\mathbf{D}_w \mathbf{B} \odot \mathbf{X}_0\|_F^2$, can be written as a function of the vectorized variable $\mathbf{X}_0$, which we denote as $\ell(\mathbf{X}_0)$. The constraint $\mathbf{B} \odot (\mathbf{X}_{met} + \mathbf{X}_{mm}) = \mathbf{S}$ can be rewritten as $\mathbf{A}\mathbf{X}_1 + \mathbf{A}\mathbf{X}_2 = \mathbf{X}_0$ for a certain matrix $\mathbf{A}$. Then our optimization problem in Eq. (7) can be written in the general form of

$$\begin{aligned} \min_{\mathbf{X}_1 \in \mathcal{X}_1, \mathbf{X}_2 \in \mathcal{X}_2, \mathbf{X}_0} & g_1(\mathbf{X}_1) + g_2(\mathbf{X}_2) + \ell(\mathbf{X}_0) \\ \text{s.t.} & \mathbf{A}\mathbf{X}_1 + \mathbf{A}\mathbf{X}_2 = \mathbf{X}_0. \end{aligned} \quad (15)$$

Recognizing that this is a special case of the sharing problem in [53], we apply the convergence result provided in this work. The detailed proof is provided in the supplementary materials.

VI. DISCUSSION

We have successfully combined the physics-based data acquisition model with learned low-dimensional models for effective metabolite and macromolecule separation. Our unique strategy of combining supervised and unsupervised learning to discover component-specific low-dimensional manifolds is a novel attempt motivated by the nature of spectroscopic data. The proposed reconstruction represents a rigorous approach to leverage deep learning to solve this long-standing challenge. In contrast to the existing methods that learn end-to-end mappings, the proposed method allows the use of a general (k, t) -

space sampling operator with high flexibility in the choices of sampling designs and SNR levels, and the ability to account for B_0 inhomogeneity. While significant noise reduction can be observed in the separated metabolite and MM spectra due to the inherent denoising capability of the low-dimensional models used, it should be noted that the proposed method is focused on addressing the signal separation problem and not a substitute for spectral quantification. Our hypothesis is that a better separation will lead to improved spectral quantification of different components of interest, which has been supported by both simulation and experimental results. Meanwhile, we expect that the proposed method can be readily integrated with other more sophisticated parametric models for data generation and representation learning (both metabolites and MMs) as well as advanced quantification strategies for the separated signal components. This is beyond the scope of this work but will be investigated in future research.

While in this work we have only considered nine metabolites that are the common molecules of interest in most brain MRSI studies, especially at 3T, the proposed model learning and reconstruction methodologies are not limited by the number of metabolites considered. Adding more metabolites into the model will increase the model order to achieve the same approximation accuracy but will not cause substantially higher computation burden. For the other ^1H metabolites, it will be very challenging to quantify them reliably given the SNR level and ignoring them has a minimal bias (due to their weak signals), hence we did not include them. But more metabolites can be considered when we adapt our method for data with higher SNRs from higher field strengths (e.g., 7T). One additional thing to note is that the learned MM-specific model should be able to capture potential residual lipids spectrally overlapping with the MM peaks (if sufficient lipid removal can be achieved), because the lineshape and frequency variations introduced when generating the MM training data. Thus, small lipid residuals will not affect the model's representation capability or metabolite separation.

One important issue with the proposed method is the choice of regularization parameters. In our current implementation for practical data, we first performed a single voxel separation for the selection of λ_1 and λ_2 using the parameter values from the phantom studies. The third parameter λ_3 was then initialized based on the simulation studies and adjusted according to the discrepancy principle. Some minor fine-tuning together with λ_1 and λ_2 was performed using visual inspection of the separation reconstruction to balance SNR improvement and smoothing effects (the relative ratios between the three parameters remained the same during this step). More sophisticated parameter selection strategies can be explored in future research. Moreover, the formulation can readily be extended to incorporate other spatio-spectral constraints for improved reconstruction.

The current fully-connected network-based DAE structures and the way of handling complex-valued data may be limited in scalability. We have investigated convolutional structures (with reduced numbers of training parameters) for both FID and spectral data and found that they were not as effective as our current DAEs in terms of dimensionality reduction.

Various combinations of fully-connected and convolutional feature extraction layers, as well as choices of activation functions, are currently being explored. The current training data generation processes used relatively simplified spectral parameter distributions. While producing strong performance, this strategy does not fully exploit the information available from experimental ^1H spectroscopy data. Estimation of more sophisticated distributions using such data will be studied in future work, e.g., using kernel density estimation [54].

Although Cartesian k-space sampling has been used to demonstrate the utility of the proposed method, other sampling trajectories can be considered by generalizing the forward encoding operator without having to retrain the models (another unique advantage of our approach). The most computationally expensive step in the current algorithm is solving Eq. (8) and Eq. (9) that involves backpropagation. This is, however, a highly parallelizable process that can significantly benefit from translating the current implementation to parallel computing platforms.

VII. CONCLUSION

We have presented a new method to reconstruct and separate metabolite and MM signals for short-TE ^1H -MRSI by learning the two signal components' distinct nonlinear low-dimensional models and using the learned models as priors for reconstruction. The models were learned using two deep autoencoder based neural networks to accurately capture metabolite and MM-specific low-dimensional manifolds of their high-dimensional spectral variations. A constrained spatio-spectral reconstruction formulation that exploits the learned models for signal separation was proposed and solved by an efficient ADMM-based algorithm. Significantly improved separation over the standard parametric fitting approach has been demonstrated using both simulated and experimental short-TE brain ^1H -MRSI data. Theoretical analysis of the proposed formulation and algorithm was also provided.

REFERENCES

- [1] M. Wilson *et al.*, "Magnetic resonance spectroscopy metabolite profiles predict survival in paediatric brain tumours," *Eur. J. Cancer*, vol. 49, pp. 457–464, 2013.
- [2] P. R. Luyten *et al.*, "Metabolic imaging of patients with intracranial tumors: H-1 MR spectroscopic imaging and PET," *Radiology*, vol. 176, pp. 791–799, 1990.
- [3] J. E. Davison *et al.*, "MR spectroscopy-based brain metabolite profiling in propionic acidemia: metabolic changes in the basal ganglia during acute decompensation and effect of liver transplantation," *Orphanet J. Rare Dis.*, vol. 6, p. 19, 2011.
- [4] M. Colla *et al.*, "MR spectroscopy in Alzheimer's disease: gender differences in probabilistic learning capacity," *Neurobiol. Aging*, vol. 24, pp. 545–552, 2003.
- [5] L. Su, A. Blamire, R. Watson, J. He, L. Hayes, and J. T. O'Brien, "Whole-brain patterns of ^1H -magnetic resonance spectroscopy imaging in Alzheimer's disease and dementia with Lewy bodies," *Transl. Psychiatry*, vol. 6, p. e877, 2016.
- [6] V. Mlynárik, G. Gambarota, H. Frenkel, and R. Gruetter, "Localized short-echo-time proton MR spectroscopy with full signal-intensity acquisition," *Magn. Reson. Med.*, vol. 56, pp. 965–970, 2006.
- [7] T. W. J. Scheenen, D. W. J. Klomp, J. P. Wijnjen, and A. Heerschap, "Short echo time ^1H -MRSI of the human brain at 3T with minimal chemical shift displacement errors using adiabatic refocusing pulses," *Magn. Reson. Med.*, vol. 59, pp. 1–6, 2008.

- [8] A. Henning, A. Fuchs, J. B. Murdoch, and P. Boesiger, "Slice-selective FID acquisition, localized by outer volume suppression (FIDLOVS) for 1H-MRSI of the human brain at 7 T with minimal signal loss," *NMR Biomed.*, vol. 22, pp. 683–696, 2009.
- [9] W. Bogner, S. Gruber, S. Trattnig, and M. Chmelik, "High-resolution mapping of human brain metabolites by free induction decay 1H MRSI at 7 T," *NMR in Biomed.*, vol. 25, pp. 873–882, 2012.
- [10] C. Gasparovic et al., "Test-retest reliability and reproducibility of short-echo-time spectroscopic imaging of human brain at 3T," *Magn. Reson. Med.*, vol. 66, pp. 324–332, 2011.
- [11] C. Cudalbu, V. Mlynárik, and R. Gruetter, "Handling macromolecule signals in the quantification of the neurochemical profile," *J. Alzheimers. Dis.*, vol. 31, pp. S101–S115, 2012.
- [12] J. Penner and R. Bartha, "Semi-LASER 1H MR spectroscopy at 7 Tesla in human brain: Metabolite quantification incorporating subject-specific macromolecule removal," *Magn. Reson. Med.*, vol. 74, pp. 4–12, 2015.
- [13] M. Považan et al., "Mapping of brain macromolecules and their use for spectral processing of 1H-MRSI data with an ultra-short acquisition delay at 7 T," *Neuroimage*, vol. 121, pp. 126–135, 2015.
- [14] R. Birch, A. C. Peet, H. Dehghani, and M. Wilson, "Influence of macromolecule baseline on 1H MR spectroscopic imaging reproducibility," *Magn. Reson. Med.*, vol. 77, pp. 34–43, 2017.
- [15] U. Seeger, U. Klose, I. Mader, W. Grodd, and T. Nagele, "Parameterized evaluation of macromolecules and lipids in proton MR spectroscopy of brain diseases," *Magn. Reson. Med.*, vol. 49, pp. 19–28, 2003.
- [16] G. D. Graham, J.-H. Hwang, D. L. Rothman, and J. W. Prichard, "Spectroscopic assessment of alterations in macromolecule and small-molecule metabolites in human brain after stroke," *Stroke*, vol. 32, pp. 2797–2802, 2001.
- [17] M. Craveiro, V. Clement-Schatlo, D. Marino, R. Gruetter, and C. Cudalbu, "In vivo brain macromolecule signals in healthy and glioblastoma mouse models: 1H magnetic resonance spectroscopy, post-processing and metabolite quantification at 14.1 T," *J. Neurochem.*, vol. 129, pp. 806–815, 2014.
- [18] N. Kunz, C. Cudalbu, V. Mlynarik, P. S. Hüppi, S. Sizonenko, and R. Gruetter, "Diffusion-weighted spectroscopy: a novel approach to determine macromolecule resonances in short-echo time 1H-MRS," *Magn. Reson. Med.*, vol. 64, 2010.
- [19] K. L. Behar, D. L. Rothman, D. D. Spencer, and O. A. Petroff, "Analysis of macromolecule resonances in 1H NMR spectra of human brain," *Magn. Reson. Med.*, vol. 32, pp. 294–302, 1994.
- [20] I. Mader, U. Seeger, J. Karitzky, M. Erb, F. Schick, and U. Klose, "Proton magnetic resonance spectroscopy with metabolite nulling reveals regional differences of macromolecules in normal human brain," *J. Magn. Reson. Imag.*, vol. 16, pp. 538–546, 2002.
- [21] K. Snoussi et al., "Comparison of brain gray and white matter macromolecule resonances at 3 and 7 Tesla," *Magn. Reson. Med.*, vol. 74, pp. 607–613, 2015.
- [22] I. Giapitzakis, R. Kreis, and A. Henning, "Characterization of the macromolecular baseline with a metabolite-cycled double-inversion recovery sequence in the human brain at 9.4 T," *Proc. of Annual Conference of ISMRM*, 2016.
- [23] M. Gottschalk, I. Tropès, L. Lamalle, S. Grand, J.-F. Le Bas, and C. Segebarth, "Refined modelling of the short-T2 signal component and ensuing detection of glutamate and glutamine in short-TE, localised, 1H MR spectra of human glioma measured at 3 T," *NMR Biomed.*, vol. 29, pp. 943–951, 2016.
- [24] S. W. Provencher, "Estimation of metabolite concentrations from localized in vivo proton NMR spectra," *Magn. Reson. Med.*, vol. 30, pp. 672–679, 1993.
- [25] S. W. Provencher, "Automatic quantitation of localized in vivo 1H spectra with LCModel," *NMR Biomed.*, vol. 14, pp. 260–264, 2001.
- [26] K. Young, B. J. Soher, and A. A. Maudsley, "Automated spectral analysis II: Application of wavelet shrinkage for characterization of non-parameterized signals," *Magn. Reson. Med.*, vol. 40, pp. 816–821, 1998.
- [27] H. Ratiney, Y. Coenradie, S. Cavassila, D. van Ormondt, and D. Graveron-Demilly, "Time-domain quantitation of 1H short echo-time signals: background accommodation," *Magn. Reson. Mater. Phys.*, vol. 16, pp. 284–296, 2004.
- [28] Y. Zhang and J. Shen, "Smoothness of in vivo spectral baseline determined by mean-square error," *Magn. Reson. Med.*, vol. 72, pp. 913–922, 2014.
- [29] H. H. Lee and H. Kim, "Parameterization of spectral baseline directly from short echo time full spectra in 1H-MRS," *Magn. Reson. Med.*, vol. 78, pp. 836–847, 2017.
- [30] F. Lam, Y. Li, B. Clifford, and Z.-P. Liang, "Macromolecule mapping of the brain using ultrashort-TE acquisition and reference-based metabolite removal," *Magn. Reson. Med.*, vol. 79, pp. 2460–2469, 2018.
- [31] C. Ma, F. Lam, Q. Ning, C. L. Johnson, and Z.-P. Liang, "High-resolution 1H-MRSI of the brain using short-TE SPICE," *Magn. Reson. Medicine*, vol. 77, pp. 467–479, 2017.
- [32] G. Peyré, "Manifold models for signals and images," *Comput. Vis. Image. Underst.*, vol. 113, pp. 249–260, 2009.
- [33] G. E. Hinton and R. R. Salakhutdinov, "Reducing the dimensionality of data with neural networks," *Science*, vol. 313, pp. 504–507, 2006.
- [34] J. Masci, U. Meier, D. Cireşan, and J. Schmidhuber, "Stacked convolutional auto-encoders for hierarchical feature extraction," in *Artificial Neural Networks and Machine Learning – ICANN 2011*, T. Honkela, W. Duch, M. Girolami, and S. Kaski, Eds., 2011, pp. 52–59.
- [35] J. Bengio, A. Courville, and P. Vincent, "Representation learning: A review and new perspectives," *IEEE Trans. Pattern. Anal. Mach. Intell.*, vol. 35, pp. 1798–1828, 2013.
- [36] S. S. Gurbani, S. Sheriff, A. A. Maudsley, H. Shim, and L. A. Cooper, "Incorporation of a spectral model in a convolutional neural network for accelerated spectral fitting," *Magn. Reson. Med.*, vol. 81, pp. 3346–3357, 2019.
- [37] H. H. Lee and H. Kim, "Intact metabolite spectrum mining by deep learning in proton magnetic resonance spectroscopy of the brain," *Magn. Reson. Med.*, vol. 82, pp. 33–48, 2019.
- [38] H. H. Lee and H. Kim, "Deep learning-based target metabolite isolation and big data-driven measurement uncertainty estimation in proton magnetic resonance spectroscopy of the brain," *Magn. Reson. Med.*, 2020, in Press.
- [39] N. Hatami, M. Sdika, and H. Ratiney, "Magnetic resonance spectroscopy quantification using deep learning," in *Proc. of Med. Image. Comput. Comput. Assist. Interv.*, 2018, pp. 467–475.
- [40] F. Lam, Y. Li, and X. Peng, "Constrained magnetic resonance spectroscopic imaging by learning nonlinear low-dimensional models," *IEEE Trans. Med. Imag.*, vol. 39, pp. 545–555, 2019.
- [41] R. Eslami and M. Jacob, "Robust reconstruction of MRSI data using a sparse spectral model and high resolution MRI priors," *IEEE Trans. Med. Imag.*, vol. 29, pp. 1297–1309, 2010.
- [42] L. Vanhamme, A. van den Boogaart, and S. V. Huffel, "Improved method for accurate and efficient quantification of MRS data with use of prior knowledge," *J. Magn. Reson.*, vol. 129, pp. 35–43, 1997.
- [43] F. Lam, Y. Li, R. Guo, B. Clifford, and Z.-P. Liang, "Ultrafast magnetic resonance spectroscopic imaging using SPICE with learned subspaces," *Magn. Reson. Med.*, vol. 83, pp. 377–390, 2020.
- [44] X. Qu et al., "Accelerated nuclear magnetic resonance spectroscopy with deep learning," *Angew. Chem. Int. Ed.*, vol. 59, pp. 10 297–10 300, 2020.
- [45] J. P. Haldar, D. Hernando, S.-K. Song, and Z.-P. Liang, "Anatomically constrained reconstruction from noisy data," *Magn. Reson. Med.*, vol. 59, pp. 810–818, 2008.
- [46] J. Yang and Y. Zhang, "Alternating direction algorithms for ℓ_1 -problems in compressive sensing," *SIAM J. Sci. Comput.*, vol. 33, pp. 250–278, 2011.
- [47] Y. Li, F. Lam, B. Clifford, and Z.-P. Liang, "A subspace approach to spectral quantification for MR spectroscopic imaging," *IEEE Trans. Biomed. Eng.*, vol. 64, pp. 2486–2489, 2017.
- [48] D. Stefan et al., "Quantitation of magnetic resonance spectroscopy signals: the jMRUI software package," *Meas. Sci. Technol.*, vol. 20, p. 104035, 2009.
- [49] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [50] J. Nocedal and S. J. Wright, *Numerical Optimization*. Springer, 2006.
- [51] R. A. de Graaf, *In vivo NMR spectroscopy: Principles and techniques*. John Wiley & Sons, 2019.
- [52] C. Ma, F. Lam, C. L. Johnson, and Z.-P. Liang, "Removal of nuisance signals from limited and sparse 1H MRSI data using a union-of-subspaces model," *Magn. Reson. Med.*, vol. 75, pp. 488–497, 2016.
- [53] M. Hong, Z.-Q. Luo, and M. Razaviyayn, "Convergence analysis of alternating direction method of multipliers for a family of nonconvex problems," *SIAM J. Optimiz.*, vol. 26, pp. 337–364, 2016.
- [54] V. Vapnik and S. Mukherjee, "Support vector method for multivariate density estimation," in *Proc. of Adv. Neural. Inf. Process. Syst.*, 1999, pp. 659–665.

Supplementary Materials

Proof of Theorem 1

We first restate the convergence theorem to be proved.

Theorem 1. *There exists a constant ρ_0 such that if $\rho \geq \rho_0$, every limit point of the sequence $(\mathbf{X}^{(i)}, \mathbf{S}^{(i)})$ generated by the algorithm described in (8), (9), (10), and (11) is a stationary solution of the optimization problem (7) (i.e. a solution that satisfies the KKT condition).*

We first explain more details of the transformation outlined in the proof sketch. Denote $\mathbf{X}_1 = \text{vec}(\mathbf{X}_{met})$, $\mathbf{X}_2 = \text{vec}(\mathbf{X}_{mm})$, $\mathbf{X}_0 = \text{vec}(\mathbf{S})$. Here vec is the standard vectorization operation of a matrix, i.e. write a matrix $(M_{ij})_{1 \leq i \leq n_1, 1 \leq j \leq n_2}$ as a vector of size $n_1 n_2$. The main purpose of this vectorization is to write the linear constraint in the standard form. Denote A as a diagonal matrix whose diagonal entries are all the entries of $\text{vec}(\mathbf{B})$. Then the constraint $\mathbf{B} \odot (\mathbf{X}_{met} + \mathbf{X}_{mm}) = \mathbf{S}$ can be written as $A\mathbf{X}_1 + A\mathbf{X}_2 = \mathbf{X}_0$. To be consistent, we also need to express the objective as a function of the vectorized variables. Therefore we rewrite $\lambda_1 \sum_{n=1}^N \|\mathcal{N}_{met}(\mathbf{X}_{met}^n) - \mathbf{X}_{met}^n\|_2^2$ as a function of the vectorized variable $\mathbf{X}_1$, denoted as $g_1(\mathbf{X}_1)$. It is not hard to write down the closed form expression of g_1 , but for the proof we do not need this expression, so we skip it. Similarly, we can rewrite $\lambda_2 \sum_{n=1}^N \|\mathcal{N}_{mm}(\mathbf{X}_{mm}^n) - \mathbf{X}_{mm}^n\|_2^2$ as a function of $\mathbf{X}_2$, denoted as $g_2(\mathbf{X}_2)$. We rewrite

$$\|\mathbf{d} - \Omega\{\mathbf{F}\mathbf{X}_0\}\|_2^2 + \lambda_3 \|\mathbf{D}_w \bar{\mathbf{B}} \odot \mathbf{X}_0\|_F^2, \quad (16)$$

as a function of the vectorized variable $\mathbf{X}_0$, which we denote as $\ell(\mathbf{X}_0)$.

Then our optimization problem in Eq. (7) can be written in the general form of

$$\begin{aligned} \min_{\mathbf{X}_1 \in \mathcal{X}_1, \mathbf{X}_2 \in \mathcal{X}_2, \mathbf{X}_0} & g_1(\mathbf{X}_1) + g_2(\mathbf{X}_2) + \ell(\mathbf{X}_0) \\ \text{s.t.} & A\mathbf{X}_1 + A\mathbf{X}_2 = \mathbf{X}_0. \end{aligned} \quad (17)$$

This is a special case of the sharing problem [1]. And our algorithm (8), (9), (10), (11) is a special case of Algorithm 4 in [1], which converges to a stationary point under certain technical conditions described in Assumption C in [1]. Next, we verify these conditions.

Assumption C1 in [1] contains two sub-assumptions: the objective $\ell(\mathbf{X}_0)$ is Lipschitz smooth, and the constraint coefficient matrix A is full column rank (the original sub-assumption requires the coefficient matrix for each variable is full column rank; here the coefficient matrix for $\mathbf{X}_1$ and $\mathbf{X}_2$ are the same). The objective function

$$\ell(\mathbf{X}_0) = \|\mathbf{d} - \Omega\{\mathbf{F}\mathbf{X}_0\}\|_2^2 + \lambda_3 \|\mathbf{D}_w \bar{\mathbf{B}} \odot \mathbf{X}_0\|_F^2 \quad (18)$$

is a quadratic function of $\mathbf{X}_0$, thus Lipschitz smooth, i.e., there exists L such that $\|\nabla \ell(\mathbf{X}_0) - \nabla \ell(\tilde{\mathbf{X}}_0)\|_F \leq L \|\mathbf{X}_0 - \tilde{\mathbf{X}}_0\|_F$ for any $\mathbf{X}_0, \tilde{\mathbf{X}}_0$. In fact, L can be chosen as the maximum eigenvalue of the Hessian of $\ell(\mathbf{X}_0)$. This verifies the first sub-assumption. Recall that A is a diagonal matrix whose diagonal entries are all the entries of $\text{vec}(\mathbf{B})$. Based on the MRSI physics, each entry of $\mathbf{B}$ is a non-zero number, thus A has full column rank. This verifies the second sub-assumption.

Assumption C2 in [1] contains two sub-assumptions: the subproblems are strongly convex with modulus $\gamma_1(\rho), \gamma_2(\rho)$ and $\gamma_0(\rho)$ respectively (these quantities may depend on ρ , thus appear in the bracket), and the penalty parameter ρ satisfies $\rho > \max\{2L^2/\gamma_0(\rho), 2L^2\}$. Since the feasible set of $\mathbf{X}_i$ is bounded, the minimum eigenvalue of the Hessian of $g_i(\mathbf{X}_i)$ has a lower bound $\lambda_{\min}$. For any $\delta > 0$, if we pick $\rho > (\delta - \lambda_{\min})/\min_{i,j} |B_{i,j}|$, then the subproblem defined in (8), (9) is strongly convex with modulus at least δ . The subproblem defined in (10) is strongly convex with modulus at least ρ (since the final term is a quadratic function with the second order coefficient being $\rho/2$), i.e., $\gamma_0(\rho) \geq \rho$. Thus the first sub-assumption holds. Since $\gamma_0(\rho) > \rho$, we only need to pick $\rho > \max\{\sqrt{2}L, 2L^2\}$ to satisfy the second sub-assumption $\rho > \max\{2L^2/\gamma_0(\rho), 2L^2\}$.

Assumption C3 requires that the objective function is lower bounded. In our problem, the objective function is non-negative, thus this assumption holds.

Assumption C4 requires that g_i 's are smooth functions (possibly non-convex) and are Lipschitz smooth, i.e., there exists L_i such that

$$\|\nabla g_i(\mathbf{X}_i) - \nabla g_i(\tilde{\mathbf{X}}_i)\| \leq L_i \|\mathbf{X}_i - \tilde{\mathbf{X}}_i\|, \forall \mathbf{X}_i, \tilde{\mathbf{X}}_i \in \mathcal{X}_i. \quad (19)$$

As we are using smooth activation functions in the neural networks, the neural network mappings $\mathcal{N}_{met}$ and $\mathcal{N}_{mm}$ are both smooth functions of the input. The function g_i consists of two terms, each of which is a concatenation of the quadratic loss function and a neural network mapping, thus g_i is smooth. Since the feasible set $\mathcal{X}_0$ is a bounded set, the maximum eigenvalue of $\nabla^2 g_i$ has an upper bound L_i in the feasible set. Thus there exists L_i such that (19) holds.

Now we have verified Assumption C in [1]. Applying [1, Theorem 3.4], we obtain the desired convergence result.

Numerical Phantom Generation

Segmented brain compartments and tissue fraction maps, i.e., gray matter (GM), white matter (WM), and cerebrospinal fluid (CSF), were first obtained from an experimentally acquired brain T_1 -weighted MPRAGE image (by SPM12, <https://www.fil.ion.ucl.ac.uk/spm/>). Compartmental concentrations and T_2^* /linewidths for the metabolites and MMs were assigned by adapting literature values [2–4]. More specifically, for metabolites (NAA, Glu, Cr1, Cr2, mI, Gln, Cho, GABA, GSH, and Lac), the concentrations (c_m) in the GM were 12.25, 9.25, 7.5, 4.25, 6.5, 9.25, 2.5, 1.5, 2.25, and 0.6 mM, respectively. And the T_2^* s were 250, 80, 150, 150, 150, 80, 200, 80, 70, and 220 ms, respectively. The concentrations for the WM were decreased by 30% with the T_2^* randomly decreased by 15%-20% for different molecules. In addition, the concentrations for all molecules in the CSF were assigned to be 0.1% of the concentrations in GM. For MMs (peaks located at 0.9, 1.21, 1.38, 1.63, 2.01, 2.09, 2.25, 2.61, 2.96, 3.11, 3.67, 3.8, and 3.96 ppm), the relative concentrations coefficients (b_l) in the GM were 0.72, 0.28, 0.38, 0.05, 0.45, 0.36, 0.36, 0.04, 0.2, 0.11, 0.64, 0.07, 1, respectively. And the linewidths (W_l) were 21.2, 19.16, 15.9, 7.5, 29.03, 20.53, 17.89, 5.3, 14.02, 17.89, 33.52, 11.85, 37.48 Hz, respectively. Similarly, the concentrations coefficients for the WM were decreased by 30% with the linewidths remains the same for different MMs. And the concentrations for all MMs in the CSF were also assigned to be 0.1% of the concentrations in GM. Next, a global scaling factor (similar to [4], on top of the individual coefficient variations) for MM signals were introduced to reflect experimentally observed metabolite-to-MM signal ratios. The regional concentrations and linewidths were subsequently combined using the tissue fraction maps as weightings in a voxel-by-voxel fashion, generating continuously varying concentration and linewidth maps. A lesion-like feature was generated in the image with a significantly altered metabolite ratio, i.e., a factor of 3 higher Cho concentration and a factor of 3 lower concentrations for other metabolites compared to the GM. Finally, the simulated metabolite basis ($\{v_m\}$) and the spectral parameters at different voxels were fed into Eq. (2) to synthesize FIDs at individual voxels. To mimic practical scenarios, we also simulated voxel-dependent random frequency shifts for different molecules (with mean zero and a standard deviation of 5 Hz) as well as residual B_0 field inhomogeneity (mean zero and a standard deviation of 10 Hz). Noisy data were simulated by adding complex white Gaussian noise. The SNR is defined as the ratio between the maximum NAA peak amplitude (across the NAA map) and noise standard deviation (σ):

$$\text{SNR} = \frac{\max_{\mathbf{r}} |\rho(\mathbf{r}, f_{\text{NAA}})|}{\sigma},$$

where f_{NAA} denotes the NAA peak frequency at around 2 ppm.

Supplementary Figures

This section shows the supplementary results mentioned in the main text.




Figure S1: Reconstructed and separated spectra produced by the proposed method from the same data used in Figs. 3-7 with the real parts of the spectra shown (phase corrected). The voxel locations are marked (red squares) in the anatomical images on the left. The first row shows the results from the simulation phantom, while the second and third rows show spectra from the in vivo FID-MRSI and sLASER data, respectively.




Figure S2-1: Experimental results from another volunteer (data acquired using the FID-MRSI sequence with $TR/TE = 1500/4$ ms, $FOV = 220 \times 220$ mm², slice thickness = 10 mm, matrix size = 24×24 , spectral bandwidth = 2000 Hz, and 512 FID samples): a) and b) Maps of NAA, Cr, and MM estimated from the separated signals produced by the proposed method (a) and the direct parametric fitting method (b). The maps were overlaid on an anatomical image for the matched slice; c) Spatially-resolved spectra (voxel location marked by the blue symbol) with the first and second rows showing the results from the proposed method (orange) and parametric fitting (blue), respectively. The original spectra are shown in black. A similar comparison to other data can be observed, with the results from the proposed method showing less spatial artifacts and less spectral under/over estimation.




Figure S2-2: Same as Fig. S2-1 with spectra shown in real parts.




Figure S3: Simulation results from data with a higher SNR (SNR=60) a) A representative spectrum (real parts) from the phantom illustrating the noise level; b) Mapping results from the data: maps of NAA, Cr, Cho, and MM (~ 0.9 ppm peak) from the ground truth (Gold Standard, column 1), the direct parametric fitting method (column 2), and the proposed method (column 3) are compared, along with relative ℓ_2 errors for the separated metabolite and MM signals (shown in images). With the higher SNR, reduced spatial variances and lower errors can be observed, as expected.




Figure S4: The DAE structure used. $\mathbf{X}$ denotes the collection of FID training data (each with a length T). The real and imaginary parts of these data are concatenated (thus a dimension of $2T$) to be used as input to the network. The DAE uses a fully connected “bottleneck” structure with an encoder (red) and a decoder (blue), and a middle feature layer with a dimension L (also referred to as the model order for our learned representations).




Figure S5: Experimental results from the in vivo FID-MRSI (top panel) and sLASER data (bottom panel). The top rows in each panel contain the molecular maps (NAA, Cr, Cho, Glu and MM) estimated from the separated signals produced by the proposed method (a and c), while the second rows in each panel the direct parametric fitting results (b and d). As can be seen, all the maps from the proposed method exhibit substantially less artifacts than those from the direct fitting, consistent with the comparison shown in the main text.




Figure S6: Comparison of the in-house fitting implementations against that provided by the jMRUI package. Different metabolite and MM spectra were simulated using experimental spectral parameters. As can be seen, our implementation produced almost the same results as those from jMRUI.

References

- [1] M. Hong, Z.-Q. Luo, and M. Razaviyayn, “Convergence analysis of alternating direction method of multipliers for a family of nonconvex problems,” *SIAM J. Optimz.*, vol. 26, pp. 337–364, 2016.
- [2] R. A. de Graaf, *In vivo NMR spectroscopy: Principles and techniques*. John Wiley & Sons, 2019.
- [3] R. Birch, A. C. Peet, H. Dehghani, and M. Wilson, “Influence of macromolecule baseline on ^1H MR spectroscopic imaging reproducibility,” *Magn. Reson. Med.*, vol. 77, pp. 34–43, 2017.
- [4] H. H. Lee and H. Kim, “Intact metabolite spectrum mining by deep learning in proton magnetic resonance spectroscopy of the brain,” *Magn. Reson. Med.*, vol. 82, pp. 33 – 48, 2019.