



Accelerated Analysis in MR Spectroscopic Imaging

Dhritiman Das

Vollständiger Abdruck der von der Fakultät für Informatik der Technischen Universität München zur Erlangung des akademischen Grades eines

Doktors der Naturwissenschaften (Dr. rer. nat.)

genehmigten Dissertation.

Vorsitzender:

Prof. Dr. Stephan Jonas

Prüfende der Dissertation:

1. Prof. Dr. Björn H. Menze
2. Priv.-Doz. Dr. Marion I. Menzel

Die Dissertation wurde am 24.02.2020 bei der Technischen Universität München eingereicht und durch die Fakultät für Informatik am 22.05.2020 angenommen.



Abstract

MR Spectroscopic Imaging (MRSI) is an advanced imaging method that is based on the principle of nuclear magnetic resonance and supports the non-invasive measurement of the in vivo metabolite levels in the tissues to be examined by generating metabolic maps. These maps show the concentration and distribution of the metabolites in the tissue sample to be examined and their precise quantification is important for the diagnosis of diseases. This modality has proven to be highly sensitive and specific in the detection of tumors and other neurological disorders. Despite its advantages, the widespread clinical use of MRSI was extremely limited due to computational bottlenecks in acquisition, processing and data analysis. In this thesis, we develop and implement data-driven methods to meet these challenges and to accelerate and improve the analysis of MR spectra data. First we develop a new *denoising* method based on frequency-phase non-local means (NLM) to improve the spectral signal-to-noise ratio (SNR). Next, we develop a machine learning architecture to quantify metabolites in the presence of noise and artifacts while significantly reducing computation time. We validate all of our methods on synthetic and human in vivo brain MRS data and demonstrate the general improvements made to the processing and analysis pipeline in MRSI.



Zusammenfassung

MR Spectroscopic Imaging (MRSI) ist eine fortschrittliche bildgebende Methode, die auf dem Prinzip der Kernspinresonanz basiert und die nichtinvasive Messung der in vivo-Metabolitenspiegel in den zu untersuchenden Geweben durch Generierung von Stoffwechselkarten unterstützt. Diese Karten zeigen die Konzentration und Verteilung der Metaboliten in der zu untersuchenden Gewebeprobe und deren genaue Quantifizierung ist wichtig für die Diagnose von Krankheiten. Diese Modalität hat sich beim Nachweis von Tumoren und anderen neurologischen Störungen als hochempfindlich und spezifisch erwiesen. Trotz seiner Vorteile war der weit verbreitete klinische Einsatz von MRSI aufgrund von rechnerischen Engpässen bei der Erfassung, Verarbeitung und Datenanalyse äußerst begrenzt. In dieser Arbeit entwickeln und implementieren wir datengetriebene Methoden, um diesen Herausforderungen zu begegnen und die Analyse von MR-Spektrendaten zu beschleunigen und zu verbessern. Zuerst entwickeln wir eine neue *denoising* Methode, die auf [frequency-phase non-local means \(NLM\)](#) basiert, um das spektrale [signal-to-noise ratio \(SNR\)](#) zu verbessern. Als Nächstes entwickeln wir eine Architektur für maschinelles Lernen, um Metaboliten in Gegenwart von Rauschen und Artefakten zu quantifizieren und gleichzeitig die Rechenzeit erheblich zu verkürzen. Wir validieren alle unsere Methoden auf synthetischen und humanen In-vivo-Hirn-MRS-Daten und demonstrieren die allgemeinen Verbesserungen, die an der Verarbeitungs- und Analyse-Pipeline in MRSI vorgenommen wurden.



Acknowledgements

I entered the Ph.D program at TU Munich assuming that the life of a doctoral student would offer an idyllic setting for me to learn, and expand my knowledge, about machine learning and medical image analysis, read all I want, have great discussions with smarter colleagues over some chilled weißbier, publish a paper or two to justify receiving my paycheck and to be able to do all this while soaking in the beautiful Bavarian countryside. Looking back, life has turned out to more hectic and intense than this, but it has been an exceptional ride!

I have had an enjoyable experience here largely because of the people I became acquainted with over the last 4 years. First of all, I would like to thank my advisor, Prof. Dr. Bjoern H. Menze, who offered me a position in his group about four years ago and played a crucial role in my decision to relocate to Munich. Starting my doctoral research and being completely unfamiliar with MR Spectroscopy, Bjoern helped me in navigating through most of the initial challenges while settling in here. He has tried to instill in me the ability to write and think clearly and provided insights into the finer nuances of the field while encouraging me to push the envelope. I am grateful for his practical advice and guidance, sharing his personal experiences and being patient with me and my unorganized ways, as well as for his constant support throughout the ups and downs that make up a Ph.D. Also, my sincere gratitude to my co-supervisor, Prof. Dr. Axel Haase, for his guidance and for demonstrating how to combine scientific expertise and leadership skills to manage the research environment at TU Munich. A special mention goes to my thesis reviewer, PD. Dr. Marion Menzel, and to Prof. Dr. Stephan Jonas for acting as the chairman and examiner for my thesis defense at a very short notice.

Next, I would like to acknowledge the support of my colleagues at GE, especially my manager, Timo Schirmer, and my co-PI and mentor, Rolf

Schulte. The GE Healthcare (formerly Global Research) group under Timo's management was an energetic and enriching place to be. The opportunity to work in the healthcare industry allowed me to better understand the practical implications of my research and to see if it can transition to a clinical and research workflow. Rolf has been a great mentor and I have enjoyed working with him. I thank him for the constant support and encouragement he has given me right from my first day at GE and TUM. A lot was made easy by having easy access to a huge pool of MRS data and the in-house analysis tools - most of which were developed by Rolf and his students. I also want to thank all my colleagues at GE - Eduardo, Ming, Xin, Pedro, Miguel, Jonathan, Tobias, Fateh, Beatriz, Anabea, Florian, Andre, Guido, Marion and Cristina. You all have been tremendously kind and supportive - I will remember the conversations we have had about the latest trends in MR and other miscellaneous things over lunch, ISMRM visits to the farthest corners of Earth and the competitive rounds of kicker, and, most importantly, I will miss the wonderful Christmas dinners at the Freising Volkshockshule that I was fortunate to have been a part of.

My doctoral position was funded by a generous Marie-Curie fellowship from the MacSeNeT EU-ITN. I would like to thank Helen Cooper and Mark Plumbley from the University of Surrey for managing the network and organizing different workshops all across Europe. The training events have helped shape my researcher skills and, now in the last sprint of my program, I am truly reaping the benefits from the intense training programs that I attended. A shout out to all my colleagues of the MacSeNeT and SpARtAN ITNs who made attending these events all the more fun - your infectious enthusiasm made this phase even more memorable.

I would like to acknowledge the great research collaborations that have made a lot of this work possible. Thanks to Mike Davies and Ian Marshall for hosting me for almost 4 months at Edinburgh in their research team and being open to my new ideas. This was an engaging and productive stint and helped in taking forward my work on machine-learning in MRS. Also, Jason Crane and Yan Li from the University of California, San Francisco for collaborating on some of the newer deep-learning work and providing such a large volume of glioma dataset to further the experimental studies.

My colleagues and friends in Munich, particularly Abhijit Guha Roy, Merkouris Simos, Anjany Sekuboyina, Suprosanna Shit, Markus Rempfler, Cagdas Ulas, Judith Zimmermann, Carolin Pirkl, Giles Tetteh, Diana Waldmannstetter, Yu Zhao, Hongwei Li, Lina Xu, Amir Bayat, Ivan Ezhov,

Sebastian Endt and John LaMaster. They contributed to an enjoyable work environment and are responsible for many good memories of Munich - every conversation, conference visit, weekend travel, pop-culture discussion and culinary exploration has been a great experience. Also, the support staff at the Graduate School of Bioengineering (Anja Drescher and the BERTI team) and the Department of Informatik have been very helpful and made it easier for me to navigate through the red-tape and other administrative hurdles.

Beyond this, my friends from far away- Allan and Pavan for being good friends since our Bachelor studies and for all random-turned-insightful chats we have had whenever time permitted. A big thanks to Shadab Khan for being a great friend and mentor as well as for being free with his time and ideas, and for sharing his experience as a researcher (and as a photographer) to help me become better. This, and the amazing hiking trips, photography sessions and drives across uncharted landscapes while clowning around - I look forward to having more such amazing adventures.

To my family - thanks to my parents for unconditionally supporting me and giving me the freedom and opportunity to pursue my dreams. My sister, Somja, and my cousins, Smritika and Shreemant, for the phone-calls, the holiday visits and the never-ending buffoonery which has been a constant source of entertainment and a major stress-reliever. It is only after soaking in these moments that one is able to get back to work with renewed energy and a clear outlook. I look forward to seeing you move higher in life and being successful in your endeavors. This one is for many more wonderful memories in the coming years.

Last, but not least, I would like to deeply thank my grandparents for their never-ending love and support ever since I was a kid, and for providing a comforting and nurturing environment at home. Words are not enough to thank them so I will not try.

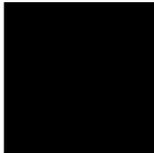


Contents

Abstract	i
Zusammenfassung	iii
Acknowledgements	v
Contents	ix
List of Figures	xiii
List of Tables	xvii
Acronyms	xix
1 Introduction	1
1.1 Challenges in MRSI	2
1.2 Summary of Contributions	6
1.3 Organization	9
2 Background	11
2.1 Magnetic Resonance Spectroscopic Imaging (MRSI)	11
2.1.1 Nuclear Magnetic Resonance Spectroscopy	11
2.1.2 MRSI	12
2.2 Acceleration in MRSI	14
2.2.1 Accelerated acquisition	14
2.2.2 Data-driven methods	14
2.3 Machine Learning	19

3	Adaptive Spectral Denoising for MR Spectroscopic Imaging using Frequency-Phase Non-Local Means	23
3.1	Introduction	24
3.2	Methods	25
3.2.1	Non-Local Means (NLM) in Frequency-Phase space	26
3.2.2	Spectral Dephasing	27
3.2.3	Spectral Rephasing and Recombination	28
3.3	Experiments and Results	28
3.3.1	Data Acquisition	28
3.3.2	Results	29
3.4	Conclusion	31
4	Quantification of Metabolites in Magnetic Resonance Spectroscopic Imaging using Machine Learning	33
4.1	Introduction	35
4.2	Methods	37
4.3	Experiments and Results	38
4.3.1	Data	38
4.3.2	Results	40
4.4	Conclusion	44
5	Direct Estimation of Model Parameters in MR Spectroscopic Imaging using Deep Neural Networks	45
5.1	Abstract	46
5.2	Introduction	46
5.3	Methods	46
5.3.1	Results	48
5.4	Discussion and Conclusion	50
6	Concluding Remarks	53
	Appendices	59
A	Metabolite Quantification of ^1H MRSI spectra in Multiple Sclerosis: A Machine Learning Approach	59
A.1	Abstract	59
A.2	Purpose	59

A.3	Methods	60
A.4	Results	62
A.5	Discussion	62
B	Using Deep Neural Networks for Metabolite Quantification in Healthy and Glioma spectra	65
B.1	Abstract	65
B.2	Architecture design	65
B.3	Data and Experiments	66
B.4	Results and Discussion	66
C	List of Publications	69
	Peer-reviewed Conference Proceedings	69
	Journal Articles	69
	Peer-reviewed Conference Abstract Proceedings	70
D	Original Publications	71
	Bibliography	97



List of Figures

- 1.1 Illustration an example healthy, human brain MRSI dataset. (A) The simulated brain with the region of interest (red box). (B) Highlighted regions corresponding to GM, WM and CSF (c) Corresponding spectrum of GM, WM and CSF. Note that CSF has only water. 2
- 1.2 Example of the different challenges in MR Spectroscopic Imaging (MRSI) acquisition and analysis such low signal-to-noise ratio (SNR), poor resolution and high noise levels. 3
- 1.3 An example of the spectral fitting performed by the LCModel [7]. The GUI displays the accuracy of fitting for a given spectra (the fit being displayed in red) while a table is co-displayed indicating the concentration estimates and fitting errors for the individual metabolites. 5

- 3.1 MRSI data dephasing shown here for a sample voxel: (A) Changes in spectral pattern as it is shifted by different phase angles. (B) Corresponding 2D frequency-phase image space generated for the voxel. A sample patch (black box) is selected and then denoised by the NLM-based matching in the frequency-phase space. . . . 27
- 3.2 Simulated brain MRSI dataset. (A) The simulated brain with the region of interest (red box). (B) Highlighted regions corresponding to GM, WM and CSF (c) Corresponding spectrum of GM, WM and CSF. Note that CSF has only water. 29

LIST OF FIGURES

3.3 NLM Denoising results in the simulated data. (From Top) Row 1 (L-R): Full SNR of NAA in – original data, with additive noise 2σ , and noise 3σ (σ is the standard deviation of the original data). Row 2 & 3: 32×32 Region of Interest (ROI) for applying the frequency-phase NLM: SNR of NAA in the original data, noise level 2σ , 3σ and the corresponding spectra of reference WM voxel (red box). Row 4 & 5 (L-R): Denoised SNR for noise level 2σ (SNR improvement = 2.9), for noise level 3σ (SNR improvement = 2.2), and the corresponding denoised spectrum. The SNR improves significantly while retaining the spatial-spectral resolution (seen by no voxel bleeding in the CSF). 30

3.4 Denoising results for in-vivo data. (A) Original human in-vivo brain MRSI data with the excitation region shown (white grid). (From Top) Row 1 & 2: SNR of NAA and the corresponding WM voxel spectra in: (B) 6-averages data (**ground-truth**), (C) 3-averages data (original) and (D) denoised, (E) single-scan data (original) and (F) denoised with corresponding spectra. Row 3: LCModel based absolute concentration estimate of NAA in: 6-averages data, 3-averages data (original) and denoised, 1-average data (original) and denoised. The NAA concentration estimate and spectral SNR improve considerably as seen in columns D (SNR = 23.29) and F (SNR = 11.38) against the ground-truth (SNR = 11.44). 31

4.1 Example brain 2D MRSI dataset. (A) The simulated brain with the region of interest (red box). (B) Highlighted regions corresponding to GM, WM and CSF (c) Corresponding spectrum of GM, WM and CSF. 37

4.2 Out-Of-Bound (OOB) Error for Simulated Spectra. The experiment is performed for a varying number of features (from 1 to 256 as shown in the legend) and each iteration is assessed for a varying number of trees (as shown in the X-axis). The Y-axis represents the OOB Error rate. The error rate is minimal for more 64 features and also converges when the number of trees is close to 100. 40

4.3	Regression Scores for the following parameters (from left to right): NAA/Cr concentration estimate and Cho/Cr concentration estimate. The X-axis represents the true values of the parameter while the y-axis represents the estimated values. Both sets of values are plotted using linear regression.	41
4.4	Synthetic-Synthetic (Spec) : Estimation error for different metabolite concentration ratios. Whiskers span the [min max] values. Median error values are represented by the red line and are as follows: NAA/Cr Regression = 0.064, NAA/Cr LCMoel = 0.077, Cho/Cr Regression = 0.043, Cho/Cr LCMoel = 0.070.	41
4.5	Synthetic (Spec)-Real (MRS Images) : Estimation error for different metabolite concentration ratios. Whiskers span the [min max] values. Median error values are represented by the red line and are as follows: NAA/Cr = 0.024, Cho/Cr = 0.034.	42
4.6	NAA/Cr and Cho/Cr concentration distribution estimates from random forest regression and non-linear model fit.	42
5.1	An example spectra generated using the basis sets provided by the ISMRM MRS Fitting Challenge 2016. Using the same basis sets, over 1 million spectra are generated with variations in NAA, Cho, Cr and other metabolite concentrations along with changes in macro-molecular baseline, lipids, linewidth (t2) and SNR.	48
5.2	Estimation error for different metabolite concentration ratios in a synthetic spectra test-set. Whiskers span the [min max] values. Median error values are represented by the red line and are as follows: NAA/Cr MLP = 0.050, LCMoel = 0.065, Cho/Cr MLP = 0.0505, LCMoel = 0.050.	49
5.3	Bland-Altman plots [68] representing LCMoel and Multi-Layer Perceptron (MLP) estimates of spectra for the real CSI dataset. The X-Axis is the mean of the LCMoel and MLP estimate, while the Y-Axis represents the relative error of the MLP estimate over the LCMoel. Bland-Altman plots for: (Left) : NAA and (Right) Choline. Both plots show a good correlation with very few outliers.	50

LIST OF FIGURES

5.4	NAA/Cr and Cho/Cr concentration distribution estimates from (Left) LCModel fit and (Right) multi-layer perceptron (MLP). The mean relative errors are 0.31 for NAA/Cr and 0.12 for Cho/Cr.	51
A.1	Pipeline for Machine Learning based quantification of MS spectra. Using the K-fold cross validation approach, this method is repeated using each of the 42 patient MS-spectra as test spectra.	60
A.2	MRSI data acquired from an SPMS patient and the corresponding spectra after voxel-wise LCModel fitting [71].	62
A.3	Bland-Altman plots [68] representing LCModel and Random Forest (RF) estimates of spectra from a sample MS patient dataset. The X-Axis is the mean of the LCModel and RF estimate, while the Y-Axis represents the relative error of the RF estimate over the LCModel. From (L-R) : Bland-Altman Plots for (Top row)-NAA, Choline and Creatine, (Bottom row)- myo-Inositol and Glx and spectra plots of outlier spectrum with high quantification error and one of the spectra with low quantification error (these are marked in the Bland-Altman plots within a red triangle and blue square respectively).	63
A.4	Absolute Relative Error box-plots for the entire 42-patient spectra dataset. The mean relative errors (red line) are: Left NAA (5.4%), Choline (10.4%), Creatine (10.5%), Right mI (15.7%), Glx (13.8%). 8 patients had an average error of > 15% per metabolite which contributed to the slight increase in the mean error across all the patients.	64
A.5	Absolute Relative Error box-plots for the sample patient spectra. The mean relative errors (red line) are: (Left) NAA (13.4%), Choline (14.1%), Creatine (11.5%), (Right) mI (11.5%), Glx (10.8%).	64
B.1	Example of long-TE glioma dataset from a sample patient. . . .	67
B.2	Metabolite maps following quantification using 3 different architectures. The MLP gives the better distribution out of all 3 while the CNN comes close though it ends up underestimating the metabolite concentration distribution for NAA.	67



List of Tables

- 3.1 SNR and LCM Quantification results for NAA, Cho and Cr before and after using frequency-phase NLM on in-vivo MRSI data. Mean SNR for the denoised data is comparable or better than the ground-truth data while the FWHM of the water peak is lower than the ground-truth data thereby preserving spatial-spectral resolution. 32

- 4.1 Concentration-ratio estimate errors using random forest regression. Results are for the experiments **Real(spec)-Real(spec)** and **Real(spec)-Real(images)**. The errors are calculated over the respective LCModel estimates as per the formula given in Eq.A.2. The major metabolites (**NAA** and **Cr**) show a low error while the smaller concentration metabolites (**mI** and **Glx**) show a slightly higher error. 43
- 4.2 Concentration-ratio estimate errors using random forest regression. Results are for the experiments **Real(spec)-Real(spec)** and **Real(spec)-Real(Images)**. The errors are calculated over the respective LCModel estimates as per the formula given in Eq.A.2. The major metabolites (**NAA** and **Cr**) show a low error while the smaller concentration metabolites (**mI** and **Glx**) show a slightly higher error. 43

- B.1 Concentration estimate errors using random forests, multi-layer perceptrons and convolutional neural networks on 40 glioma patients. A patient-wise K-fold cross validation was performed and the results shown are averaged over all the folds. Overall, the cnn ends up out-performing the other 2 architectures in terms of metabolite quantification. 68



Acronyms

CNN	convolutional neural network
FCN	fully convolutional neural network
MLP	multi-layer perceptron
MRSI	MR Spectroscopic Imaging
NLM	frequency-phase non-local means
NMR	nuclear magnetic resonance
NN	neural network
RF	random forest
SNR	signal-to-noise ratio
SVM	support vector machine

Introduction

Magnetic Resonance Spectroscopic Imaging (MRSI), also known as chemical shift imaging (CSI) is a clinical imaging modality for studying tissues *in-vivo* to investigate and diagnose neurological diseases. More specifically, this modality can be used in non-invasive diagnosis and characterization of patho-physiological changes by measuring specific tissue metabolites in the brain. It is based on the concept of nuclear magnetic resonance and works by exploiting the resonance frequency of a molecule, which depends on its chemical structure, to obtain information about the concentration of a particular metabolite [1]. The time-domain complex signal of a nuclei is given by:

$$S(t) = \int p(\omega) \exp(-i\Phi) \exp(-t/T_2^*) d\omega \quad (1.1)$$

The frequency-domain signal is given by $S(\omega)$, T_2^* is the magnetization decay in the transverse plane due to magnetic field inhomogeneity and $p(\omega)$ comprises of Lorentzian absorption and dispersion line-shapes function having the spectroscopic information about the sample. Φ represents the phase, $(\omega t + \omega_0)$, of the acquired signal where ωt is the time-varying phase change and ω_0 is the initial phase. For the acquired MRSI data, I , Φ is unknown. This process allows generation of metabolic maps through non-linear fitting to estimate concentration of metabolites such as N-acetyl-aspartate (NAA), Creatine (Cr) and Choline (Cho).

Signal and metabolite quantitation [2] gives MRSI an increased advantage over other forms of medical magnetic resonance imaging. The majority of the other techniques produce high-resolution images with visible image contrast differences that can be identified by a trained person. This enables the detection of lesions or other anatomical abnormalities by only a visual

1. INTRODUCTION

examination. MRSI, in addition to this, also provides a visual representation of the metabolite information [3]. However, it also suffers from a number of drawbacks which makes it difficult to use MRSI as the imaging modality of choice in clinics.

Figure 1.1 shows an example of spectra in a healthy, human brain MRS dataset.

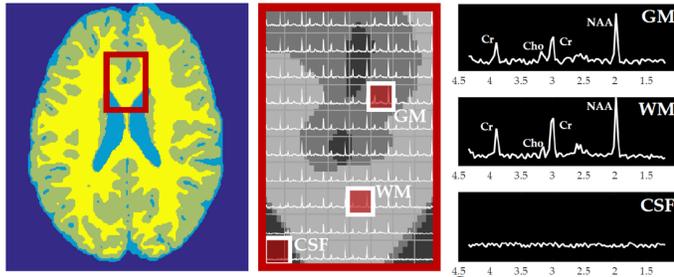


Figure 1.1: Illustration an example healthy, human brain MRSI dataset. (A) The simulated brain with the region of interest (red box). (B) Highlighted regions corresponding to GM, WM and CSF (c) Corresponding spectrum of GM, WM and CSF. Note that CSF has only water.

In the next section, we discuss the challenges in MRSI and outline the two main applications of this thesis towards accelerating the analysis and processing pipeline in MR Spectroscopic Imaging - using data driven methods such as denoising algorithms and machine learning. Subsequently, we summarize our main contributions in Section 1.2 and outline the remaining manuscript in Section 1.3.

1.1 Challenges in MRSI

Figure 1.2 shows the challenges in MRSI acquisition and analysis. These are discussed in detail in this section.

Low signal-to-noise ratio (SNR): The water signal- which are used for proton images in conventional MRI- has a concentration which is approximately 10,000 times higher than the concentration of metabolites in a typically acquired MRSI data leading to extremely low sensitivity.

Long acquisition times: Due to the low SNR challenge, multiple signal averages are often acquired to improve the SNR in MRSI experiments.

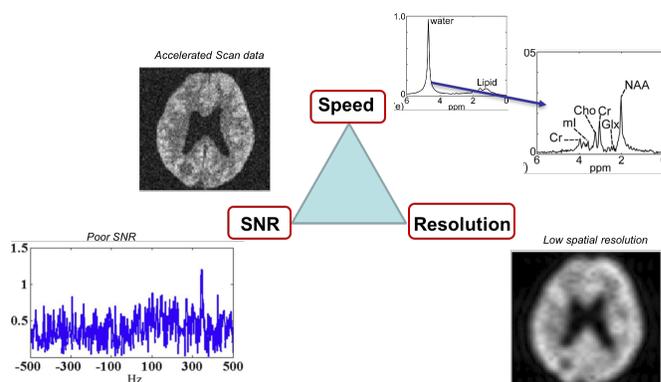


Figure 1.2: Example of the different challenges in MR Spectroscopic Imaging (MRSI) acquisition and analysis such as low signal-to-noise ratio (SNR), poor resolution and high noise levels.

Furthermore, conventional MRSI methods mostly rely on very slow spatial encoding schemes [2,15]. These approaches significantly prolong the data acquisition time and serve as major deterrents for clinicians to use this modality, especially for in vivo studies.

Poor spatial resolution and long acquisition times: Conventional MRSI methods mostly rely on very slow spatial encoding schemes which lead to long acquisition times and restrict the practical use of MRSI in clinical applications. To overcome this drawback, existing MRSI methods involve constrain the number of spatial encodings and data reconstructions with large voxel sizes (on the order of 1 cm^3) to ensure sufficient SNR. This results in poor spatial resolution and low volume coverage make it difficult to visually interpret the available information. As mentioned earlier, the water signal, having a concentration which is approximately 10,000 times higher than the concentration of metabolites, also adds to the difficulty in obtain spatially-resolved spectra [4]. One approach to deal with this is to acquire the data at long echo times (TE) as water signal decays with TE more quickly than the metabolite signals [5][6]. However, this approach ends up reducing the signal amplitude of the metabolites.

Noise-sensitive spectral fitting and metabolite quantification:

During acquisition, the signals collected from the scanner consist of field induction decays (FIDs) which are composed of several sinusoidal components oscillating at the resonance frequencies of the protons they represent. Upon completion of acquisition, these FIDs are Fourier-transformed into

1. INTRODUCTION

frequency domain signals. In order to extract meaningful information about the tissues being investigated, the next step in the MRSI pipeline is quantification of the individual metabolites present in the spectra acquired. This, in MRSI terminology, is referred to as “metabolite quantification” or “spectral fitting”. A common model for describing the shape of proton resonance in a magnetic field is the Lorentzian-Gaussian function as shown in Eq. 1.1. In the frequency domain, the parameters for this function include: frequency, amplitude, zero- and first-order phase, and two linewidth coefficients representing the combination of Lorentzian and Gaussian dampening. Algorithms which model a spectrum as a combination of individual peak resonances are referred to, appropriately, as “peak fitting” algorithms.

Mathematically, each metabolite has a corresponding “basis set of signal” developed consisting of a combination of known Lorentzian-Gaussian resonances known as “multiplets”. In the MRSI domain, and for the work in this thesis, the LCModel [7], a commercial model-fitting software tool, is considered to be the gold standard model for spectral fitting [7]. A software package developed by Stephen Provencher in 2001 that utilizes basis sets to model spectra. This tool works on the same premise of using the basis sets to iteratively fit to the acquired spectra (which serves as input) using a non-linear least-squares fit and thereby quantify most of the metabolites present in the input spectrum. However, the spectral fitting method produces high-error rates especially when metabolites have low concentration, or have overlapping resonance peaks with other metabolites or spurious signals such as lipids and macromolecules. The fitting tool also exhibits a high-error rate if the spectrum has broadened peaks or a high amount of noise leading to local minima [8]. An example of a typical LCModel fit has been shown in 1.3.

There are also other algorithms being used with varying degrees of success. These include the QUEST [9] method or “quantitation based on semi-parametric quantum estimation” which operates on time-domain signals and applies singular value decomposition (SVD) to separate the metabolite signals from the background signals and retain the important signal components, including modeling the smooth components of the macromolecular baselines. On the other hand, the TARQUIN [10] algorithm (Totally Automatic Robust Quantitation in NMR) involves suppressing the water signal in the acquired spectra the water signal using the Hankel SVD followed by a non-linear least squares fitting for quantifying the metabolites. The MIDAS software (Metabolite Imaging and Data Analysis System) has the FITT tool

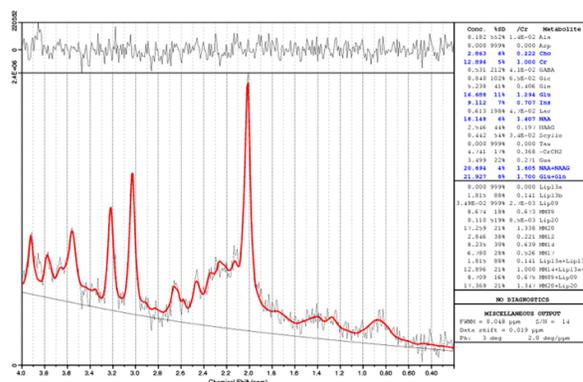


Figure 1.3: An example of the spectral fitting performed by the LCModel [7]. The GUI displays the accuracy of fitting for a given spectra (the fit being displayed in red) while a table is co-displayed indicating the concentration estimates and fitting errors for the individual metabolites.

for spectral fitting which operates on frequency-domain spectra acquired during MRSI scans using the accelerated EPSI sequence. It, again as the other models, uses the Lorentzian-Gaussian model for fitting the the peak components of an acquired spectrum, and smoothing splines for modeling the macromolecular baseline.

A common thread running through the above fitting tools is the use of common theme is the need for iterative fitting methods, such as non-linear least squares or orthogonal matching pursuit among others, to decompose the input spectra into its constituent metabolites through signal-peak extraction. These algorithms involve a significant computational burden: the LCModel (considered as a gold standard) can take upto 1 hour for fitting a whole-brain MRSI data (as an example, we consider a short-TE scan having a size of 32x32 voxels with 256 spectral points). Additionally, the parameters need to be manually-tuned for each dataset (depending on acquisition parameters and complexity of data) leading to issues in parallelization and scaling. This, in addition, to the high-error rates due to high noise and peak-broadening constitute a serious constraint on the practical implementation of this technology in the research and clinical environments.

One way of improving the SNR of the signal is by using post-processing methods such as denoising algorithms [11], apodization (Gaussian/Lorentzian), filter-based smoothing and transform-based methods [12]. However, these methods reduce resolution and remove important quantifiable information by

averaging out the lower-concentration metabolites. Recently, data-dependent approaches such as the Non-Local Means (NLM), which use the redundancy inherent in periodic images, are being used extensively for denoising [12] [13]. Other similar methods incorporate methods from machine learning and pattern recognition (such as sub-space estimation, super-resolution, dictionary learning) to improve the SNR and resolution of the MRSI data.

Spectral fitting, on the other hand, can be considerably improved by adopting machine-learning methods for eliminating low-quality spectra and performing metabolite quantification with robustness to noise and other artifacts.

1.2 Summary of Contributions

This thesis is set in the context of the analysis and processing pipeline in MR Spectroscopic Imaging. As highlighted in the previous section, we have selected two major challenges in MRSI processing that, when addressed, can significantly improve the processing pipeline for MRSI data and ensure a successful transition towards the clinical use of this modality: First, the low signal-to-noise ratio (SNR) of the MR spectra and the need for repeated acquisitions for signal avergaing, and second, the quantification of metabolites present in the spectra with a focus on improving the computational time and robustness to noise and other artifacts.

In the following, we give a brief introduction to the setting of each publication-based chapter and summarize its content and contributions.

Chapter 3: Spatially Adaptive Spectral Denoising for MR Spectroscopic Imaging using Frequency-Phase Non-local Means

MR Spectroscopy imaging protocols usually have long scanning times in order to obtain spatially resolved spectra with an optimal signal-to-noise ratio. This is mainly the due to abundance of clinically-relevant metabolites which have a concentration which is approximately 10,000 times smaller than water peak (used in conventional proton MR imaging). Furthermore, from a processing perspective, an optimal SNR is essential for spectral fitting or metabolite quantification as the non-linear voxel-wise fitting to noisy data

leads to a high amount of local minima and noise amplification resulting in poor spatial resolution[7].

Conventionally, post-processing methods such as denoising algorithms[14], apodization (Gaussian/Lorentzian) [15], filter-based smoothing [16] and transform-based methods [17] have been used to improve the SNR of MR spectra. However, these methods tend to reduce resolution and remove important quantifiable information as a result of over-smoothing the spectra and thereby reducing the signal amplitude lower-concentration metabolites. Recently, data-dependent approaches such as the Non-local Means (NLM) [13], which use the redundancy inherent in periodic images, are being used extensively for denoising [12]. In the case of MRSI, this periodicity implies that the spectra (or a metabolite region of it) in a given voxel may have similar spectra or metabolite region in other voxels of the same dataset within its frequency-phase space. Therefore, the algorithm may, hypothetically, carry out a weighted average of the voxels within this frequency-phase space, depending on the similarity of the spectral information of their neighborhoods to the neighborhood of the voxel to be denoised.

We contribute a method for spectrally adaptive denoising of MRSI spectra in the frequency-phase space based on the concept of Non-local Means that incorporates the redundant information from the same, acquired dataset thereby reducing the need for any additional prior information. To address the the lack of arbitrary phase-information in the acquired spectra we also implement a "dephasing" approach on the spectral data which increases the amount of redundant frequency-phase information within the voxels of the MRSI data. This improves the effectiveness of the NLM algorithm while a selection of modified, anisotropic patch sizes prevents over-smoothing of metabolite peaks. We include experiments using simulated brain data and healthy human in-vivo 2D MRSI datasets and show that the proposed method increases the SNR of the data while retaining the spatial resolution of the spectra. More importantly, the increase in SNR for a single-scan compared to multiple-acquisition averages offers a alternative to also accelerate MRSI imaging by acquiring fewer averages.

Chapter 4: Quantification of Metabolites in Magnetic Resonance Spectroscopic Imaging using Machine Learning

Quantification of metabolites in the acquired spectra is a crucial part of the MRSI analysis workflow. This gives the concentration of the metabolites present in the underlying tissue being investigated and obtaining accurate concentration estimates is imperative for detecting tumors and diagnosing other metabolic diseases within the tissue. Conventionally, model-fitting tools such as the LCModel are used for metabolite quantification. This software tool uses a linear combination of metabolite basis spectra set to model the spectral measurement in the frequency domain and incorporates the prior knowledge of the data while iteratively modeling the fit using a non-linear optimisation. This ensures robustness in the model leading to estimation of the spectral parameters such as concentration of metabolites. Although it is the gold-standard fitting tool for MRSI data, it is often hindered by its drawbacks such as: (1) computationally-intensive and requirement of manual parameter tuning. (2) higher-quantification error for low SNR and artifact-heavy spectra due to a high amount of local minima generated during the non-linear voxel-wise fitting.

In this chapter, we propose a machine learning-based method as an alternative to the model-fitting tool for metabolite quantification. We implement the random-forest method to perform a multi-variate regression on MR spectra and aim to estimate concentrations of the major metabolites used in MRSI for brain tissue analysis. While prior work [18] has used random forests for classification and assessment of quality of spectra, this work aims to break new ground in using random forests for improving the quantification pipeline in MRSI while addressing the constraints of long fitting times, manual parameter-tuning and lack of robustness to noise and artifacts. We also simulate an extensive, representative training set comprising of synthetic spectra with varying features using a physical model and aim to generalize this learning to human in-vivo MRSI data. Our experiments involve both single-voxel and 2D MRSI brain data from healthy, human subjects, while our results show that the learning from the synthetic data is successfully generalized to the in vivo data thereby giving the accurate metabolite concentration estimates. Moreover, these metabolite concentration predictions are similar to the LCModel but are performed in

a fraction of the time.

Chapter 5: Direct Estimation of Model Parameters in MR Spectroscopic Imaging using Deep Neural Networks

In this chapter, we build upon the work done in Chapter 4 and use a more robust architecture in the form of a multi-layer perceptron (a type of artificial neural network) to predict the metabolite concentrations from healthy human in-vivo brain MRS data. We show a better prediction of the concentrations along with robustness to noise and artifacts.

1.3 Organization

This a publication-based thesis with the following structure: Chapter 1 introduced to the topic of magnetic resonance spectroscopic imaging along with the current challenges involved with it, and summarized our contributions. Chapter 2 gives a brief summary of relevant terminology and key concepts from magnetic resonance imaging theory, data-driven methods and machine learning, which are used throughout this manuscript.

Chapter 3 to 5 are composed of three publications [19, 20, 21] in their original form. They have been published as peer-reviewed conference proceedings and abstracts, and are therefore self-contained. Each of these chapters starts with a brief summary, containing the full citation of the original publication, a short synopsis of the corresponding publications content and the thesis authors contributions. In order to improve the reader's experience, the text layout of the publications was harmonized and their bibliographies have been merged into one single bibliography at the end of this document.

Chapter 6 offers discussion and conclusions over the presented material and suggest directions for future work. Appendix A and Appendix B provide additional unpublished work which complement the work presented in Chapter 4 and 5, respectively. Finally, a complete list of publications that have been written during the time period of this doctoral thesis can be found in Appendix C.

Background

The overarching themes of this thesis include magnetic resonance imaging theory including nuclear magnetic resonance, data-driven methods and machine-learning. This chapter aims at giving a concise summary of key concepts and notation used throughout this thesis, but it is not intended to be a representative overview of the most important concepts of each field. For a more complete and in-depth discussion on [MRSI](#), please refer to [1], for denoising [12, 13] and to [22, 23, 24, 25, 26] for topics of machine learning.

2.1 Magnetic Resonance Spectroscopic Imaging (MRSI)

2.1.1 Nuclear Magnetic Resonance Spectroscopy

Most MRI experiments are only concerned with one type of spin present within the same chemical environment (for example: protons on the water molecules). In practice, however, the nuclei in a biological object are often attached to different chemical environments. For example, the nuclei in different molecules are surrounded by orbiting electrons, which produce different magnetic fields that “locally” perturb the field felt by the nuclei to various extents. This effect is known as the electron shielding effect. To further simplify this, the nuclei in different chemical environments (thus subject to different electron shielding effects) resonate at different frequencies. This is described by,

$$\omega = \gamma B_0(1 - \sigma) \tag{2.1}$$

2. BACKGROUND

where B_0 is the strength of the main magnetic field and σ is the shielding constant. This frequency dispersion gives rise to the chemical shift phenomenon and resonance peaks that constitutes the core foundation of the field of NMR spectroscopy.

Due to the existence of spins resonating at different frequencies, an intrinsic frequency dimension can be introduced to mathematically describe the signal acquired in the free induction decay (FID) period (the period after the RF pulse ends). This FID signal acquired can be modeled as:

$$s(t) = \int_{\omega_M} \rho(\omega) e^{-t/T_2(\omega)} e^{-i\omega t} d\omega, \quad (2.2)$$

where $\rho(\omega)$ defines a spectral density function characterizing the frequency distribution and $T_2(\omega)$ captures the relaxation effect. Based on different assumptions for $\rho(\omega)$, different models can then be derived from Eq. 2.2 to analyze and extract information from the FID signal in spectroscopy experiments [1]. Using NMR spectroscopy to study different chemical shift frequencies and the spin densities associated with these frequencies has far-reaching impacts in chemical and biological studies. For example, it has been applied to determine the structures of various macromolecules, to quantify metabolic signatures in biological tissue samples in order to detect, diagnose and characterize pathological diseases, and to understand the basic metabolic and physiological processes in the living bodies. However, for in-vivo experiments, spectroscopy data acquired from a large volume of interest, which is excited, is inherently limited by the lack of spatial specificity, which is crucial since the metabolic processes in the human body are spatially dependent. Therefore, the concept of [MR Spectroscopic Imaging \(MRSI\)](#) is introduced to address this problem.

2.1.2 MRSI

[MRSI](#) aims to generate spatially-resolved spectroscopic information from the imaging object through combining the spectral encoding during the FID period and spatial encoding using gradients. Since additional frequency dimensions are introduced in [MRSI](#) experiments, the measured data $s(\cdot)$ can then be modeled in a (k, t) -space as

$$s(k_m, t_q) = \int_V \int_{\omega_M} \tilde{\rho}(r, f) e^{-i2\pi f \cdot t_q} e^{-i\gamma \Delta B(r) t_{1,q}} e^{-i2\pi k_m \cdot r} df dr + \xi(k_m, t_q), \quad (2.3)$$

2.1. Magnetic Resonance Spectroscopic Imaging (MRSI)

where $\tilde{\rho}(r, f)$ is the spatio-spectral function of interest (containing both spatial and spectral variations), V denotes the excited volume of interest, ω_M the spectral bandwidth (determined by the range of chemical shift), γ the gyromagnetic ratio, $\Delta(B)r$ the B_0 field inhomogeneity distribution, t_1 the first temporal dimension in \mathbf{t} and $\xi(\cdot)$ the measurement noise (modeled as white Gaussian). When integrated over the frequency domain, Eq. 2.3 can be simplified into

$$s(k_m, t_q) = \int_V \rho(r, t_q) e^{-i\gamma\Delta B(r)t_q} e^{-i2\pi k_m \cdot r} dr + \xi(k_m, t_q), \quad (2.4)$$

where $\rho(r, t)$ is the Fourier counterpart of $\tilde{\rho}(r, f)$,

$k_m = (m_x\Delta k_x, m_y\Delta k_y, m_z\Delta k_z)$ and $t_q = q\Delta t$. m_x, m_y, m_z and q are the sample indices along different spatial and spectral dimensions.

$\Delta k_x, \Delta k_y, \Delta k_z$ and Δt are the corresponding sampling intervals. With the signal models in Eqs. 2.3 and 2.4, the imaging problem in MRSI is to recover $\tilde{\rho}(r, f)$ or $\rho(r, t)$ from the set of (k, t) -space measurement $s(k_m, t_q)$. The most common approach to produce the spatio-spectral encodings for MRSI is the chemical shift imaging (CSI) method that uses pulsed gradients to encode the phase in all spatial dimensions after each excitation-step and uses the complete period of the FID for the resulting encoding of spectra. After each excitation step, all the encoding gradients end at the same time after each excitation to ensure the same initial chemical shift state. However, due to the slow spatial encoding fashion of CSI per excitation, only very limited data can be acquired in practical clinical and research experiments. This causes significant truncation artifacts, peak overlapping, interference by lipids, macromolecular signals and, subsequently, poor spatial resolution. In order to achieve high-resolution MRSI in a short imaging time, accelerated data acquisition is needed. This has brought the need for different data-driven methods to accelerate acquisition or, on the other hand, acquire less data and accelerate the processing times.

2.2 Acceleration in MRSI

A number of data-driven methods have been proposed for accelerating MRSI acquisition, signal denoising and improving spectral fitting. Some of these methods have been reviewed here.

2.2.1 Accelerated acquisition

A number of techniques have been developed to deal with these drawbacks of MRSI and attempt to ensure its wider use among clinicians. Current acquisition techniques such as Parallel Imaging [27] and Echo-Planar Spectroscopic Imaging [28] focus on accelerated scanning times combined with advanced reconstruction techniques to improve the SNR of the data. Compressed Sensing techniques have also been proposed by [29][30][31] though computational complexity issues have prevented a wider, commercial use of this method. Despite this, further accelerated acquisitions are desirable. Signal averaging is a commonly used approach in MRSI to deal with the noise present in the data. The noise level is often similar in magnitude to the metabolite signals. Taking averages of the same signal causes the signal to increase in proportion to the number of averages while the noise increase in proportion to square root of the averages. Signal averaging may be limited in case of substantial physiological motion while acquiring the data which may lead to arbitrary variations in phase and frequency and cause less than desired improvement in SNR.

2.2.2 Data-driven methods

One of the earliest data-driven methods used for accelerating MRSI acquisition and analysis was compressed sensing. Compressed Sensing (CS) is based on undersampling the data in a sparse domain (such as wavelet or DCT or any other transform domain having a high amount of redundancy), followed by reconstructing the data to obtain a denoised output [32]. However, with MRSI data, the amount of redundancy in the spatial domain is low (given the relatively small size of the voxels). One application of compressed sensing to hyperpolarized ^{13}C flyback 3D-MRSI [30] involves using the high SNR available from this hyperpolarization technique to obtain high-resolution data. Spatial coverage and resolution is limited as only a limited number of phase encodes can be fit into the short acquisition time for hyperpolarized

MRSI. Compressed sensing resolves this issue by exploiting the sparsity in the hyperpolarized 13-C spectra by under-sampling the spectral k-space resulting in enhancement of the spatial resolution by a factor of 2. Another application of compressed sensing can be seen in a four-dimensional echo-planar based J-resolved spectroscopic imaging sequence involving two spatial and two spectral dimensions [31]. One spatially encoded dimension k_x and one temporal dimension t_2 are simultaneously acquired using the EPSI readout sequence. The remaining spatial and spectral dimensions (k_y and t_1) are incrementally collected using non-uniform sampling (NUS) and the missing data is reconstructed using compressed sensing coupled with the split Bregman method- an l_1 -norm minimization algorithm. The CS-based reconstruction also performs a data-dependent denoising to remove noise from the original under-sampled data. This approach works favorably for approximately 20 – 25% of the original data implying an acceleration factor of 4-5 times.

Wavelet based methods. Experiments on denoising MRS data were conducted which involved automatic decomposition of MRS based on Principal Component Analysis (PCA) and Independent Component Analysis (ICA) [33]. Earlier work used PCA to project MRS into a subspace which was coordinated by orthogonal principal components [34]. This method, however, failed in the case of small data size and low SNR due to difficulty in separating noise from MRS data in the orthogonal subspace. Moreover, as the ICA method is highly noise-sensitive, the independent components (ICs) extracted from small size data denoised using PCA (having complex noise features) are of poor quality. As an improvement over this method, wavelet was combined with PCA to eliminate noise of varying levels of complexity. This involved using ICA to decompose MRS data into suitable ICs having biochemical properties as corresponding features. This enabled using ICA for classification purposes and subsequent tuning of the wavelet packet parameters leading to removal of the higher-order noise from the spectra.

Low-rank approximation methods. In general, low-rank approximation methods applied to spatial-spectral data incorporate low-rank approximations [11][35] in the $k - t$ domain by assuming partial separability (PS) between spatial and temporal variations at low orders. Additionally, these methods are complemented by exploiting the low-rank of the spectral signal due to its linear predictability.

Further improvements on this method involve extending the PS model to

2. BACKGROUND

incorporate the tissue boundary constraints from the high-resolution MRSI data to enable better low-rank filtering [36]. The boundary constraints allow the classification of voxels having weak and/or localized metabolite signals. Therefore, the local low-rank filtering ensures that these signals can be protected from being smoothed or removed by global low-rank approximation. In addition to this, the method further improves the low-rank model by integrating B_0 field inhomogeneity correction of MRSI data by using the surrogate nuclear norm rank-minimization method. The method significantly improves the SNR of the MRSI spectra though it is susceptible to segmentation errors (while classifying low-rank and high-rank spectral signals) which can degrade the denoising performance.

A new method known as SPICE (SPectroscopic Imaging by exploiting spatioSpectral CorrElation) [37][38] uses the sub-space model for both data acquisition and image reconstruction while providing a better tradeoff for SNR, resolution and speed. For the data acquisition, a hybrid echo-planar spectroscopic imaging sequence for 2D MRSI is used which allows for an extended (k,t)-space coverage in a short acquisition period [39][40][41]. The B_0 inhomogeneity field effects from the acquired data are removed by utilizing an available high-resolution field map in a regularized super-resolution reconstruction scheme. After determining the sub-space, the reconstruction method then involves estimating the set of spatial coefficients (with a lower number of degrees of freedom). A low-rank model and field inhomogeneity correction method (as mentioned above) along with an edge-preserving regularization are incorporated in a linear least-squares formulation. These additional priors help in obtaining an improved reconstruction. The method achieves fast acquisition of high-resolution data along with high SNR for both phantom and in-vivo 2D and 3D MRSI data. From the acquisition side, further optimization of the pulse sequences are desired. This is because the proposed sequence is susceptible to chemical shift displacement and can have only a limited range of echo times. Better pulse sequences would include shorter TE times and reduced chemical shift displacement errors to improve spatial localization. Additionally, integration of parallel imaging to the SPICE framework can provide a better trade off between SNR, speed and resolution allowing this method to be used for practical applications. From the sampling and reconstruction side, optimal (k,t)-space sampling strategies, such as compressed-sensing based sparse sampling schemes, can be explored to improve speed and spatioSpectral encoding. The reconstruction pipeline requires a high computational load and the implementation time is quite

high. This can be improved by incorporating better spatio-spectral/temporal constraints and using hybrid data-sets for joint estimation of subspace structure and spatial coefficients. Other developments in optimization schemes could also improve computational efficiency of the method thereby allowing for a faster reconstruction.

Denoising methods. This section deals with state-of-the-art reconstruction and denoising techniques which have been applied to other imaging tools such as dynamic MRI, perfusion MRI [42] and which can possibly be adapted and applied to MRSI.

A number of methods have been proposed for better metabolite detection and fitting of MRS images [9] [43]. Using a formulation based on combining a Gaussian Markov random field (GMRF) prior with a frequency-domain model for the free induction decay, spatial smoothness of selected parameter maps are assumed as spatial priors along with commonly used prior knowledge [8]. The spatial smoothness prior can be modified for certain image-specific parameters such as phase or line-width which allows for estimation of parameters, such as amplitude of the signal and variance, which may be of diagnostic relevance. This method of using spatial prior knowledge shows considerable benefits such as low variance and improved details of the estimated parameter maps, better resolution of the overlapping peaks of the metabolites choline and creatine, and higher SNR of the data. Further improvements to this method can be made by including multimodal spatial information from morphological MR images (such as segmented tissue maps) to adjust the coupling of neighboring voxels and provide a better spatial fitting of MRSI data.

Another method focuses on improved concentration estimation of metabolites such as Glutamate (Glu) and Glutamine (Gln) [43]. Glu and Gln occur at very low concentrations with poor resolution between their peaks [43]. The LCModel [7] (the gold-standard basis fitting algorithm for MR spectra) gives a combined estimate of Glu and Gln due to its inability to resolve the two peaks. The high resolution spectral analysis method known as 'Spectral Zooming' estimates the unique power spectral density (PSD) which corresponds to the maximum entropy solution of a zero-mean stationary Gaussian process. this method then computes several PSDs of the metabolites from a moving window of the measured data to give a separate estimate of Glu and Gln. The method successfully gives separate concentration estimates for both Glu and Gln (better than the LCModel estimates) and is validated using 3 different phantom data-sets and 1 human in-vivo data. One area in

2. BACKGROUND

which this method can be improved is in accounting for gray-white volume fractions for proper estimation of the concentration at each voxel and also providing absolute concentration values.

Other denoising methods from the computer vision and/or machine learning domain, which have been applied to MRI [42] [13], offer the possibility of being integrated in the post-processing pipeline of MRS Imaging. The method of Non-local MRI upsampling [44] proposes a data-adaptive patch-based reconstruction combined with a subsampling coherence constraint to upsample the acquired data to a higher resolution. This method is partially similar to single image super-resolution (SR) techniques based on self-similarity using a regularization expression. It involves a denoising step based on an iterative reconstruction-correction scheme (constrained by non-ideal sampling) followed by a patch-based non-local reconstruction to take advantage of the local redundancy in MR images. The patch-based approach prevents excessive blurring while preserving the edge-information. The proposed method, though useful, exhibits a high computational burden and can lead to loss of information of smaller features having a contrast similar to noise levels. Another method reconstructs data using the kernel regression along the temporal manifold structure of the undersampled k-space MRI data [45]. This allows for fast reconstruction of dynamic MRI from undersampled k-space data, thus leading to accelerated acquisition. The manifold learning is facilitated by using Laplacian Eigenmaps which preserve the data structure by ensuring that the data points (in the k-space samples) which are closely located to each other in the high-dimensional space remain close to each other in the low-dimensional sub-space as well. This method achieves a significantly faster reconstruction of dynamic MRI data with higher SNR as compared to the state-of-the-art compressed sensing schemes though better sampling strategies can be chosen for increased robustness. The kernel regression-based subspace approach is similar to the SPICE method discussed earlier and can be extended to MRSI data as well.

Other patch-based approaches include the [frequency-phase non-local means \(NLM\)](#) method [12] which has been applied to MRI data as well [13]. The NLM method exploits the redundancy inherent in periodic images such as MRI slices to perform denoising. It does a full search across the data to find voxels having intensity similar to the reference voxel following which it performs a weighted averaging of all similar voxels to give a denoised reference voxel - the weights are determined by the distance between the similar voxels and the reference voxels. The method is widely used and can be suitably

extended to [MRSI](#) by incorporating MRS features such as replacing intensity comparison with spectral data comparison, frequency(ppm)-specific search and phase-shift methods (to account for arbitrary phase and inhomogeneity issues with [MRSI](#) data). The computational burden of such a method is high and can be made better by using advanced dictionary learning techniques and more powerful optimization schemes. A recent state-of-the-art method known as Filter Forests has been proposed for denoising MRI data. Filter Forests is a linear, non-iterative forest-based predictor which performs signal restoration by determining the optimal filtering kernels that can be applied to each data point [46]. The kernels (size and value) can be determined based on observation and learning the spatial or temporal context of the data. The forest performs a recursive partition of the input signal such that a simple convolution kernel can be used at each leaf after which the forest is then trained to minimize the regularized least-squares error of the kernel at the corresponding leaf. This method has been used for a variety of image processing tasks such as image denoising, depth image refinement and 1D signal magnitude estimation while achieving speed and accuracy significantly higher than the state of the art. As this data-driven method uses the spatio-temporal context of the continuous variables in the data while performing regression analysis, it may be suitable for MR Spectroscopy as well where the convolution kernel can be adapted to the spectra features and replacing the 'spatial-temporal' feature of the variables with the 'spatial-spectral' feature. As a summary, [MRSI](#) reconstruction and post-processing is a challenging problem and efficient artifact removal, denoising and accelerated acquisition are the need of the hour. If this problem is viewed from the perspective of data analysis, then the computer vision and machine learning domain in general offers methods (some of which have been discussed in this review) which can be attempted on spectral data while incorporating its specific features and physical constraints [47].

2.3 Machine Learning

Machine learning is a broad field with overarching themes involving pattern recognition, data mining, and statistical learning theory. It includes algorithms that *learn* a mapping, such as classification or regression functions, from a pool of observations, by making "*sense*" of the underlying patterns in the data. The Some of the more traditional and famous machine-learning

2. BACKGROUND

algorithms include *shallow* algorithms such as the [support vector machine \(SVM\)](#) [48], the [random forest \(RF\)](#) [49] or the [neural network \(NN\)](#) [50, 51]. As an example, [random forest \(RF\)](#) have been shown to be effective in a wide range of classification and regression problems. These comprise of a set of binary trees wherein splits are created in each tree based on a random subsets of the feature variables on which the forests are subsequently trained. Piece-wise linear regression is implemented by each tree over the input data and, after seeking for the best prediction at every node, data points are sent to the left or right branches based on feature selection by thresholding. This process continues till it reaches the end of the tree and subsequently the weighted average of the prediction from each tree is taken to give a single output estimate. The randomness in the training process encourages the trees to give independent estimates which can be combined to achieve an accurate and robust result.

Newer variants employ the concepts of deep learning which *learns* the deeper, more complex patterns in the data such as [52, 53, 54, 55, 26]. [NN](#) are systems of linear operations designed to have plasticity, or the ability to update themselves iteratively based on the output and minimizing the predictive errors, so-called because of their similarity to how the brain works (in a simplified fashion). NNs are composed of nodes called neurons which interact with each other by passing information. By combining many inter-connected neurons across multiple, "*deeper*" layers, dense non-linear functions can be approximated by the simple linear combination of signals as a result of which the network *learns* the mapping of more complex, unstructured patterns in the observations. Mathematical operations such as backpropagation are used to iteratively pass information between the neurons in different layers eventually leading to the network to converge into a steady-state. In recent years, machine learning techniques have played an important role in accelerating progress across different field, especially in medical image analysis. Applications include state-of-the-art methods for segmenting brain tumors [56], liver tumors [57] or detecting stroke lesion [58] - machine learning schemes are involved at some stage of the processing and analysis workflow and have shown significant improvement in performance over the traditional, deterministic and model-based approaches.

From the perspective of [MRSI](#), machine learning tools have been used only for specific tasks such as classification of spectra [59] and assessment of spectral quality [60]. This opens up the possibility of using the recent advances in machine learning to predict mr spectra parameters while ad-

addressing the drawbacks of conventional fitting tools such as long computation time and poor performance for data with some artifacts.

Adaptive Spectral Denoising for MR Spectroscopic Imaging using Frequency-Phase Non-Local Means

This chapter has been published as **peer-reviewed conference paper**.

© Springer International Publishing AG 2016

D. Das, E. Coello, R. F. Schulte, and B. H. Menze. “Spatially Adaptive Spectral Denoising for MR Spectroscopic Imaging using Frequency-Phase Non-Local Means.” In: *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2016*. Ed. by S. Ourselin, L. Joskowicz, M. R. Sabuncu, G. Unal, and W. Wells. Cham: Springer International Publishing, 2016, pp. 596–604. DOI: [10.1007/978-3-319-46726-9_69](https://doi.org/10.1007/978-3-319-46726-9_69)

Synopsis: This work discusses the denoising problem in [MR Spectroscopic Imaging \(MRSI\)](#). We propose a novel frequency-phase non-local means method which exploits the redundant spectra information within a dataset and use this to denoise the acquired spectra. Our method is validated on both simulated spectra and a healthy human 2D mrsi dataset showing significant improvement in the acfsnr while retaining the spatial-spectral resolution..

Contributions of thesis author: algorithm design and implementation, computational experiments and composition of manuscript.

Abstract

Magnetic resonance spectroscopic imaging (MRSI) is an imaging modality used for generating metabolic maps of the tissue in-vivo. These maps show the concentration of metabolites in the sample being investigated and their accurate quantification is important to diagnose diseases. However, the major roadblocks in accurate metabolite quantification are: low spatial resolution, long scanning times, poor signal-to-noise ratio (SNR) and the subsequent noise-sensitive non-linear model fitting. In this work, we propose a frequency-phase spectral denoising method based on the concept of non-local means (NLM) that improves the robustness of data analysis and scanning times while potentially increasing spatial resolution. We evaluate our method on simulated data sets as well as on human in-vivo MRSI data. Our denoising method improves the SNR while maintaining the spatial resolution of the spectra.

3.1 Introduction

Magnetic Resonance Spectroscopic imaging (MRSI), also known as chemical shift imaging, is a clinical imaging modality for studying tissues in-vivo to investigate and diagnose neurological diseases. More specifically, this modality can be used in non-invasive diagnosis and characterization of pathophysiological changes by measuring specific tissue metabolites in the brain. Accurate metabolite quantification is a crucial requirement for effectively using MRSI for diagnostic purposes. However, a major challenge with MRSI is the long scanning time required to obtain spatially resolved spectra due to abundance of metabolites that have a concentration which is approximately 10,000 times smaller than water. Current acquisition techniques such as Parallel Imaging and Echo-Planar Spectroscopic Imaging [28] focus on accelerated scanning times combined with reconstruction techniques to improve the SNR of the spectral signal. Despite this, further accelerated acquisitions are desirable. Furthermore, an improved SNR is needed as the non-linear voxel-wise fitting to noisy data leads to a high amount of local minima and noise amplification resulting in poor spatial resolution [8].

The SNR of the signal can be improved by post-processing methods such

as denoising algorithms [11], apodization (Gaussian/Lorentzian), filter-based smoothing and transform-based methods [12]. However, these methods reduce resolution and remove important quantifiable information by averaging out the lower-concentration metabolites. Recently, data-dependent approaches such as the Non-Local Means (NLM), which use the redundancy inherent in periodic images, are being used extensively for denoising [12]. In the case of MRS, this periodicity implies that the spectra in any voxel may have similar spectra in other voxels in the frequency-phase space. Therefore, it carries out a weighted average of the voxels in this space, depending on the similarity of the spectral information of their neighborhoods to the neighborhood of the voxel to be denoised.

Our Contribution. *In this work, we propose a method for spectrally adaptive denoising of MRSI spectra in the frequency-phase space based on the concept of Non-Local Means.* Our method compensates for the lack of phase-information in the acquired spectra by implementing a dephasing approach on the spectral data. In the next section, we introduce the experimental methods beginning with the concept of NLM in the frequency-phase space followed by the spectral dephasing and rephasing approach. Our proposed method is then validated quantitatively and qualitatively using simulated brain data and human in-vivo MRSI data sets to show the improvements in SNR and spatial-spectral resolution of MRSI data.

3.2 Methods

MR Spectroscopy. Magnetic resonance spectroscopy is based on the concept of nuclear magnetic resonance (NMR). It exploits the resonance frequency of a molecule, which depends on its chemical structure, to obtain information about the concentration of a particular metabolite [1]. The time-domain complex signal of a nuclei is given by: $S(t) = \int p(\omega) \exp(-i\Phi) \exp(-t/T_2^*) d\omega$. The frequency-domain signal is given by $S(\omega)$, T_2^* is the magnetization decay in the transverse plane due to magnetic field inhomogeneity and $p(\omega)$ comprises of Lorentzian absorption and dispersion line-shapes function having the spectroscopic information about the sample. Φ represents the phase, $(\omega t + \omega_0)$, of the acquired signal where ωt is the time-varying phase change and ω_0 is the initial phase. For the acquired MRSI data, I , Φ is unknown. This process allows generation of metabolic maps through non-linear fitting to estimate concentration of metabolites

3. ADAPTIVE SPECTRAL DENOISING FOR MR SPECTROSCOPIC IMAGING USING FREQUENCY-PHASE NON-LOCAL MEANS

such as N-acetyl-aspartate (NAA), Creatine (Cr) and Choline (Cho).

3.2.1 Non-Local Means (NLM) in Frequency-Phase space

As proposed by Buades et al. [12], the Non-Local Means (NLM) method restores the intensity of voxel x_{ij} by computing a similarity-based weighted average of all the voxels in a given image. In the following, we adapt NLM to the MRSI data: let us suppose that we have complex data, $I : \Omega^3 \mapsto \mathbb{C}$ of size $M \times N$ and noisy spectra $S_{ij}(\omega)$, where $(x_{ij} | i \in [1, M], j \in [1, N])$ and Ω^3 is the frequency-phase grid. Using NLM for denoising, the restored spectra $\hat{S}_{ij}(\omega)$ is computed as the weighted average of all other spectra in the frequency-phase space defined as:

$$\hat{S}_{ij}(\omega) = \sum_{x_{kl} \in \Omega^3} w(x_{ij}, x_{kl}) S_{kl}(\omega) \quad (3.1)$$

As a probabilistic interpretation, spectral data $S_{11}(\omega), \dots, S_{MN}(\omega)$ of voxels x_{11}, \dots, x_{MN} respectively are considered as MN random variables X_{ij} and the weighted average estimate $\hat{S}_{ij}(\omega)$ is the maximum likelihood estimate of $S_{ij}(\omega)$.

$N_{ij} = (2p + 1)^3$, $p \in \mathbb{N}$ is the cubic neighborhood of voxel x_{ij} within the search volume $V_{ij} = (2R + 1)^3$ around x_{ij} along the frequency, phase and spatial directions. $R \in \mathbb{N}$, where R is the radius of search centered at the voxel x_{ij} . The weight $w(x_{ij}, x_{kl})$ serves as a quantifiable similarity metric between the neighborhoods N_{ij} and N_{kl} of the voxels x_{ij} and x_{kl} provided $w(x_{ij}, x_{kl}) \in [0, 1]$ and $\sum w(x_{ij}, x_{kl}) = 1$. The Gaussian-weighted Euclidean distance is computed between $S(N_{ij})$ and $S(N_{kl})$ as shown below:

$$w(x_{ij}, x_{kl}) = \frac{1}{Z_{ij}} e^{-\frac{\|S(N_{ij}) - S(N_{kl})\|_2^2}{h^2}} \quad (3.2)$$

where Z_{ij} serves as the normalization constant such that $\sum_j w(x_{ij}, x_{kl}) = 1$, $S(N_{ij})$ and $S(N_{kl})$ are vectors containing the spectra of neighborhoods N_{ij} and N_{kl} of voxels x_{ij} and x_{kl} respectively and h serves as a smoothing parameter [13].

To increase the robustness of our method for MRSI data, in the next section we propose a dephasing approach tailored for use in the frequency-phase NLM.

3.2.2 Spectral Dephasing

For the acquired data I , as the spectral phase $\Phi(I)$ is unknown, the probability of finding a similar neighborhood spectra are very low. To counter this effect, a dephasing step is performed to consider a wide range of possible phase variations in the pattern analysis. For each voxel x_{ij} , the complex time-domain signal $S_{ij}(t)$ is shifted by a set of phase angles Θ . This is given by $S_{ij}^{\Theta}(t) = S_{ij}(t).e^{-i\Theta}$, where $\Theta \in [-n_1\pi, (n_2 + 2)\pi]$, $\{n_1, n_2 \in \mathbb{R} \mid n_1, n_2 \geq 0\}$. Θ here is defined to be the range of angles through which the spectrum can be shifted. The dephased signal is transformed into the frequency-domain, $S_{ij}^{\Theta}(t) \xrightarrow{\mathcal{F}} S_{ij}^{\Theta}(\omega)$, following which its real component, $\text{R}(S_{ij}^{\Theta}(\omega))$, is taken to generate a 2D spectral-phase matrix. Note that in this 2D matrix generated, for each voxel x_{ij} , the imaginary part at a given Θ is $\text{I}(\Theta) = \text{R}(\Theta + \pi/2)$. This approach is illustrated in Fig. 3.1.

Repeating this step for all MN voxels gives us a 3-D dataset on which the NLM is implemented to give the denoised spectra $\hat{S}_{ij}^{\Theta}(\omega) \in \mathbb{R}$. Our approach has 2 key innovations: the denoising method is (i) robust to phase shifts as the range of angles considered varies from 0 to 2π periodically for all spectral signals, and is (ii) adaptive to the imaging sequence as the spectrum is denoised by relying on similar signals in the given data and not on predefined prior assumptions.

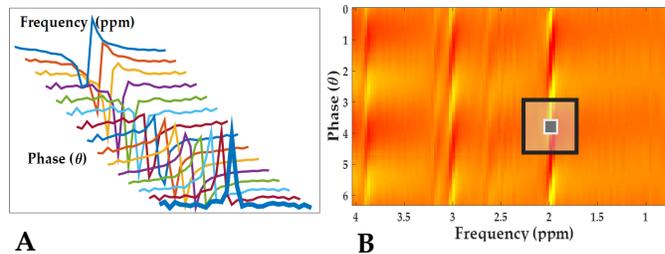


Figure 3.1: MRSI data dephasing shown here for a sample voxel: (A) Changes in spectral pattern as it is shifted by different phase angles. (B) Corresponding 2D frequency-phase image space generated for the voxel. A sample patch (black box) is selected and then denoised by the NLM-based matching in the frequency-phase space.

3.2.3 Spectral Rephasing and Recombination

Post-NLM, $\hat{S}_{ij}^{\Theta}(\omega)$ is rephased in order to generate the denoised complex signal $C(\hat{S}_{ij}^{\Theta}(\omega))$. The complex spectral signal $C(\hat{S}_{ij}^{\Theta}(\omega))$ is re-generated $\forall \Theta$ by combining $R(\hat{S}_{ij}^{\Theta}(\omega))$ and $I(\hat{S}_{ij}^{\Theta}(\omega))$ ($= R(\hat{S}_{ij}^{\Theta+\pi/2}(\omega))$). The equivalent time signal is obtained by $C(\hat{S}_{ij}^{\Theta}(\omega)) \xrightarrow{\mathcal{F}^{-1}} C(\hat{S}_{ij}^{\Theta}(t))$. After this, $C(\hat{S}_{ij}^{\Theta}(t))$ undergoes an inverse phase shift by $-\Theta$ to remove the dephasing effect as given by $C(\hat{S}_{ij}^{-\Theta}(t)) = C(\hat{S}_{ij}^{\Theta}(t)).e^{i\Theta}$. This re-phased signal is transformed back to the spectral domain to obtain $C(\hat{S}_{ij}^{-\Theta}(\omega))$. Thereafter, the $C(\hat{S}_{ij}^{-\Theta}(\omega))$ are averaged over all Θ to generate a single complex spectra $C(\hat{S}_{ij}(\omega))$.

3.3 Experiments and Results

We performed two different experiments to test the improvement in SNR and metabolite quantification using our proposed denoising method. In the first experiment, we evaluate our method on the publicly available BrainWeb database [61], while in the second experiment we use human in-vivo MRSI data. The SNR of a metabolite was calculated by dividing the maximum value of the metabolite peak by the standard deviation of the spectral region having pure noise. For both experiments, we tested with different noise levels against a ground-truth data to assess the improvement in SNR and spatial-spectral resolution.

3.3.1 Data Acquisition

We used BrainWeb to simulate a brain MRSI image (size: 64x64 voxels, slice thickness = 1mm, noise level = 3%) with segmented tissue types, namely White Matter (WM), Grey Matter (GM) and Cerebro Spinal Fluid (CSF) as shown in Fig. 4.1. In order to have a comparable spectrum with the in-vivo data, water was added to the signal and the main metabolites-NAA, Cho and Cr- were simulated using Priorset (Vespa) [62]. Metabolite concentrations for WM, GM and CSF were based on commonly reported literature values [2] [63]. Next, Gaussian noise of levels 2 and 3 times the standard deviation, σ , of the original image data were added to the ground-truth signal.

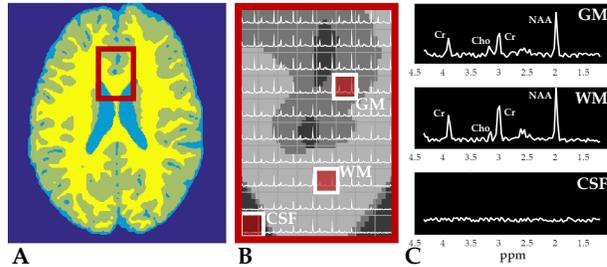


Figure 3.2: Simulated brain MRSI dataset. (A) The simulated brain with the region of interest (red box). (B) Highlighted regions corresponding to GM, WM and CSF (c) Corresponding spectrum of GM, WM and CSF. Note that CSF has only water.

In the case of in-vivo data, we acquired a 2D-MRSI data of the brain of a healthy human volunteer using a 3T-HDxt system (GE-Healthcare). PRESS localization [64], CHESSE water suppression [65] and EPSI read-out [28] were used as part of the sequence. The acquisition parameters were: Field of view (FOV) = 160x160x10 mm³, voxel size = 10x10x10 mm³, TE/TR = 35/2000 ms and spectral bandwidth = 1 kHz. The dataset was zero-filled and reconstructed to generate a grid of 32x32 voxels and 256 spectral points. 6 (**ground truth**), 3 and 1 averages were acquired with a total scan duration of 33 minutes (5.5 minutes per average). Fig. 3.4 (A) shows the in-vivo data acquired along with the entire field-of-view (white grid) and the corresponding spectra of a voxel (red box).

3.3.2 Results

Simulated data: In Fig. 3.3, we show the SNR improvement for NAA for data with noise levels 2σ and 3σ in a 32×32 region of interest. It is evident that while the spectral SNR improves significantly, the spatial resolution is preserved as the lower concentration metabolite peaks have only a small amount of smoothing and there is no voxel bleeding in the CSF (containing only water).

In-vivo data: Fig. 3.4 reports the SNR improvement in NAA for the 3-averages and the 1-average data as compared to the ground-truth 6-averages data. The figure also presents the results from the LCModel [7] which is the gold standard quantitation tool in MRS analysis. LCModel fits the spectral signal $S(\omega)$ using a basis set of spectra of metabolites acquired

3. ADAPTIVE SPECTRAL DENOISING FOR MR SPECTROSCOPIC IMAGING USING FREQUENCY-PHASE NON-LOCAL MEANS

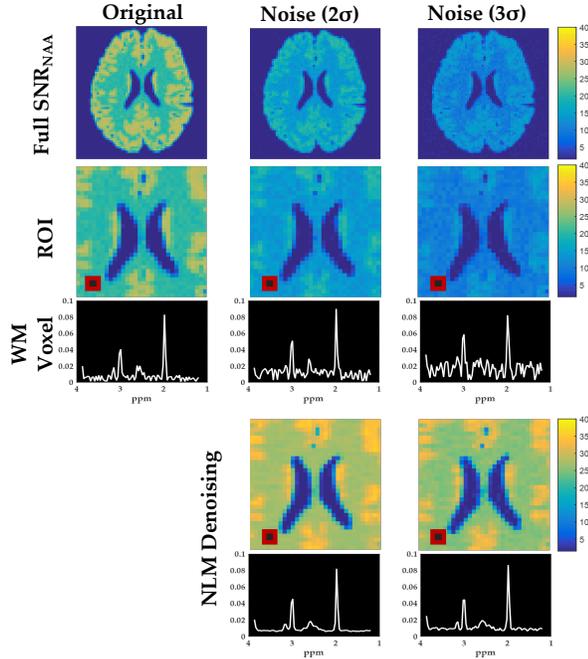


Figure 3.3: NLM Denoising results in the simulated data. (From Top) Row 1 (L-R): Full SNR of NAA in – original data, with additive noise 2σ , and noise 3σ (σ is the standard deviation of the original data). Row 2 & 3: 32×32 Region of Interest (ROI) for applying the frequency-phase NLM: SNR of NAA in the original data, noise level 2σ , 3σ and the corresponding spectra of reference WM voxel (red box). Row 4 & 5 (L-R): Denoised SNR for noise level 2σ (SNR improvement = 2.9), for noise level 3σ (SNR improvement = 2.2), and the corresponding denoised spectrum. The SNR improves significantly while retaining the spatial-spectral resolution (seen by no voxel bleeding in the CSF).

under identical acquisition conditions as the in-vivo data. As explained earlier for noisy data, the non-linear fitting leads to poor spatial resolution. Therefore, the LCModel can be used to assess the improvement in spatial resolution through a better fit. Due to space constraints, we present the results for NAA only and mention the SNR values for Cho and Cr. LCModel quantification (Fig. 3.4) shows that the absolute concentration estimation of NAA in the denoised data improves significantly. The Full-Width Half Maximum (FWHM) shows information about the water peak – a narrow

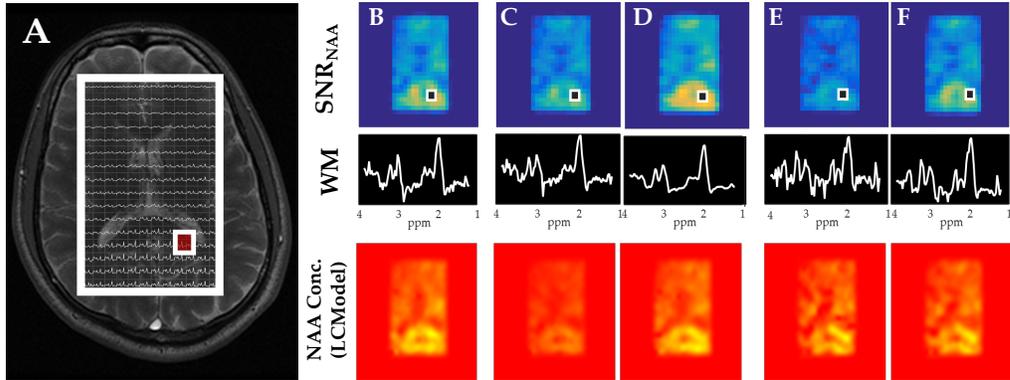


Figure 3.4: Denoising results for in-vivo data. (A) Original human in-vivo brain MRSI data with the excitation region shown (white grid). (From Top) Row 1 & 2: SNR of NAA and the corresponding WM voxel spectra in: (B) 6-averages data (**ground-truth**), (C) 3-averages data (original) and (D) denoised, (E) single-scan data (original) and (F) denoised with corresponding spectra. Row 3: LCMoDel based absolute concentration estimate of NAA in: 6-averages data, 3-averages data (original) and denoised, 1-average data (original) and denoised. The NAA concentration estimate and spectral SNR improve considerably as seen in columns D (SNR = 23.29) and F (SNR = 11.38) against the ground-truth (SNR = 11.44).

peak gives a better spatial resolution. As shown in Table 3.1, the FWHM of the denoised 1- and 3-averages data is lower than the 6-averages data while the corresponding mean SNR improves considerably. Therefore, we observe here that our method can accelerate MRSI data acquisition by almost 2 times by reducing the number of scans acquired.

3.4 Conclusion

In this work, we proposed a novel frequency-phase NLM-based denoising method for MRS Imaging to improve the SNR and spatial resolution of the metabolites. A spectral dephasing approach is promoted to compensate for the unknown phase information of the acquired data. To the best of our knowledge, this is a novel application of the concept of NLM and has been validated on both simulated and in-vivo MRSI data.

3. ADAPTIVE SPECTRAL DENOISING FOR MR SPECTROSCOPIC IMAGING USING FREQUENCY-PHASE NON-LOCAL MEANS

Table 3.1: SNR and LCM Quantification results for NAA, Cho and Cr before and after using frequency-phase NLM on in-vivo MRSI data. Mean SNR for the denoised data is comparable or better than the ground-truth data while the FWHM of the water peak is lower than the ground-truth data thereby preserving spatial-spectral resolution.

Data	Mean SNR (NAA)	SNR improvement (NAA)	Mean SNR (Cho)	SNR improvement (Cho)	Mean SNR (Cr)	SNR improvement (Cr)	FWHM
1-average	5.70	-	3.74	-	4.04	-	0.125
3-averages	8.82	-	5.59	-	5.97	-	0.135
6-averages (ground-truth)	11.44	-	6.96	-	7.61	-	0.135
Single scan (NLM)	11.38	1.98	6.89	1.82	7.78	1.90	0.130
3-averages (NLM)	23.29	2.63	14.08	2.48	15.59	2.57	0.134

In particular, we assessed the effect of our method on metabolites such as NAA, Cho and Cr and obtained a visible improvement in SNR while the spatial resolution was preserved which, subsequently, led to a better estimation of the absolute concentration distribution of NAA. This has direct benefits as it would accelerate data acquisition by taking fewer scan averages. Future work would involve using a more robust metabolite-specific search in the given dataset with less smoothing. This can be coupled with optimal computational efficiency and better estimation of the in-vivo metabolites.

Acknowledgements. The research leading to these results has received funding from the European Union’s H2020 Framework Programme (H2020-MSCA-ITN-2014) under grant agreement n° 642685 MacSeNet.

Quantification of Metabolites in Magnetic Resonance Spectroscopic Imaging using Machine Learning

This chapter has been published as **peer-reviewed conference paper**.

© Springer International Publishing AG 2017

D. Das, E. Coello, R. F. Schulte, and B. H. Menze. “Quantification of Metabolites in Magnetic Resonance Spectroscopic Imaging Using Machine Learning.” In: *Medical Image Computing and Computer Assisted Intervention, MICCAI 2017*. Ed. by M. Descoteaux, L. Maier-Hein, A. Franz, P. Jannin, D. L. Collins, and S. Duchesne. Cham: Springer International Publishing, 2017, pp. 462–470. DOI: [10.1007/978-3-319-66179-7_53](https://doi.org/10.1007/978-3-319-66179-7_53)

Synopsis: This work deals with the metabolite quantification problem in [MR Spectroscopic Imaging \(MRSI\)](#). We propose a machine-learning method using random-forests for estimating metabolite concentrations. Additionally, we also create a pipeline for simulating synthetic training spectra having similar features as short-TE human in-vivo brain spectra. Finally, we evaluate our method on single-voxel and [MRSI](#) healthy human brain spectra and benchmark these against the LCMoDel to evaluate improvement in speed and prediction error..

Contributions of thesis author: algorithm design and implementation,

4. QUANTIFICATION OF METABOLITES IN MAGNETIC RESONANCE SPECTROSCOPIC IMAGING USING MACHINE LEARNING

computational experiments and composition of manuscript.

Abstract

Magnetic Resonance Spectroscopic Imaging (MRSI) is a clinical imaging modality for measuring tissue metabolite levels in-vivo. An accurate estimation of spectral data parameters allows for better assessment of spectral quality and metabolite concentration levels. The current gold standard quantification method is the LC Model. However, this fails for spectra having poor signal-to-noise ratio (SNR) or a large number of artifacts. This paper introduces a framework based on random forest regression for accurate estimation of the output parameters of a model based analysis of MR spectroscopy data. The goal of our proposed framework is to learn the spectral features from a training set comprising of different variations of both simulated and in-vivo brain spectra and then use this learning for the subsequent metabolite quantification. Experiments involve training and testing on simulated and in-vivo human brain spectra. We estimate parameters such as concentration of metabolites and compare our results with that from the LCModel.

4.1 Introduction

Magnetic resonance spectroscopic imaging (MRSI) is an in-vivo clinical imaging modality which detects nuclear magnetic resonance signals produced by nuclei in living tissues. Quantification of this signal amplitude generates metabolic maps which show the concentration of metabolites in the sample being investigated. Accurate quantification of these metabolites is important for diagnosis of brain tumor and other in-vivo diseases. For this purpose, a common practice in the MRS community has been to use non-linear spectral fitting tools such as the LCModel [7], TARQUIN [10], AMARES [9] and ProFit amongst which the LCModel is regarded as the gold standard fitting tool. In this study, we present an alternative to the non-linear model fitting using a machine learning approach.

Non-linear model fitting. The LCModel software uses a linear combination of metabolite basis spectra set to model the spectral measurement in the frequency domain. It also uses smoothing splines to model the baseline

4. QUANTIFICATION OF METABOLITES IN MAGNETIC RESONANCE SPECTROSCOPIC IMAGING USING MACHINE LEARNING

signals and subsequently fits the parameters of the basis set using a non-linear optimisation. LCModel incorporates the prior knowledge of the data while modeling the fit and this ensures robustness in the model leading to estimation of the spectral parameters such as concentration of metabolites. Some of the drawbacks of this non-linear fitting model are: (1) Metabolite quantification can be time-consuming depending on the dataset size and requires a lot of manual parameter tuning. (2) The error in estimating parameters is lower if high SNR spectra is used since the non-linear voxel-wise fitting to noisy data leads to a high amount of local minima and subsequent inaccuracy in quantification. [59][8].

Machine Learning. Machine learning methods such as decision forests, random forests [66] are being extensively used in the medical imaging community for various tasks such as parameter estimation, diseases diagnostics, segmentation, etc. In MRSI, machine learning tools have been used only for specific tasks such as classification of spectra [59] and assessment of spectral quality [60]. This opens up the possibility of using the recent advances in machine learning to predict MRSI data parameters while addressing the drawbacks of conventional fitting tools such as long computation time and poor performance for data with some artifacts.

Our Contribution. In this work, we propose a simple yet effective method using random forest regression for multi-parameter estimation in MR Spectroscopic Imaging. We generate over 1 million simulated spectra training-set having concentration magnitudes, linewidth effects, baseline and lipid artifacts. We also use spectral data from 287 human subjects to create a physical training model to be used in the regression framework (Sect. 4.3.1). In the following we present our method adapting random forest regression to MRSI (Sect. 4.2) followed by experiments in the aforementioned dataset. Our proposed method is then validated quantitatively and qualitatively using: (1) synthetic brain spectra, (2) human in-vivo single voxel spectra having the same image acquisition protocol as the physical training model and (3) independently acquired human in-vivo 2D MRS Images to perform a blind test on the physical and synthetic models. We present the results (Sect. 4.3.2) of our experiments followed by a summary and discussion (Sect. 5.4) on the future work in this domain. This is the first application- to the best of our knowledge- of machine learning for determining MRS parameters which were otherwise determined using basis fitting tools.

4.2 Methods

MR Spectroscopy. Magnetic resonance spectroscopy, based on the concept of nuclear magnetic resonance (NMR), exploits the resonance frequency of a molecule, to obtain information about the concentration of a particular metabolite [1]. The time-domain complex signal of a nuclei is given by:

$$S(t) = \int p(\omega)\exp(-i\Phi)\exp(-t/T_2^*)d\omega. \quad (4.1)$$

The frequency-domain signal is given by $S(\omega)$, T_2^* is the magnetization decay in the transverse plane due to magnetic field inhomogeneity and $p(\omega)$ comprises of Lorentzian absorption and dispersion line-shapes function having the spectroscopic information about the sample. Φ represents the phase, $(\omega t + \omega_0)$, of the acquired signal where ωt is the time-varying phase change and ω_0 is the initial phase. Non-linear fitting tools facilitate the generation of metabolic maps to estimate concentration of metabolites such as N-acetyl-aspartate (NAA), Creatine (Cr) and Choline (Cho). An example of the spectra present in the brain has been shown in Fig. 4.1.

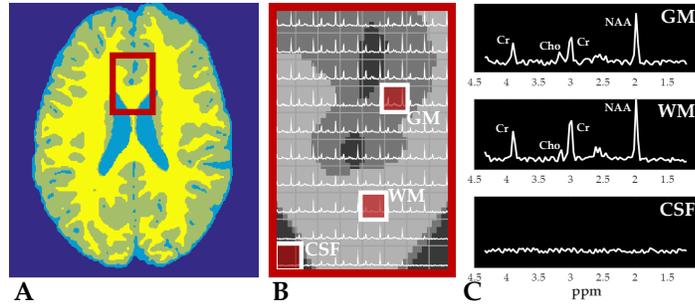


Figure 4.1: Example brain 2D MRSI dataset. (A) The simulated brain with the region of interest (red box). (B) Highlighted regions corresponding to GM, WM and CSF (c) Corresponding spectrum of GM, WM and CSF.

Random Forest Regression. Random Forests [66] have been shown to be effective in a wide range of classification and regression problems. These comprise of a set of binary trees wherein splits are created in each tree based on a random subsets of the feature variables on which the forests are subsequently trained. Piecewise linear regression is implemented by each tree over the input data and, after seeking for the best prediction at every

4. QUANTIFICATION OF METABOLITES IN MAGNETIC RESONANCE SPECTROSCOPIC IMAGING USING MACHINE LEARNING

node, data points are sent to the left or right branches based on feature selection by thresholding. This process continues till it reaches the end of the tree and subsequently the weighted average of the prediction from each tree is taken to give a single output estimate. The randomness in the training process encourages the trees to give independent estimates which can be combined to achieve an accurate and robust result.

For MRSI, we adapt the random forest approach to have a training dataset $D = (S_i, Y_i)$, $i \in [1, N]$, where N is the total number of training spectra. S_i represents the training spectral data while Y_i represents the corresponding multi-parameter training labels. For our model, we consider the concentrations of NAA, Cho and Cr for simulated data, while for the real data we additionally consider Myo-Inositol (mI) and Glutamate+Glutamine (Glx). Therefore, for a given spectra S_i , $Y_i = [NAA_i, Cho_i, Cr_i, mI_i, Glx_i]$.

Running the random forest regression on this produces a training model which can then be used to obtain parameter estimates \hat{Y}_j of test spectra S_j having test labels Y_j , $j \in [1, M]$ where M is the total number of test spectra.

Error Calculation. For our experiments, given the estimate \hat{Y}_j and the testing label Y_j , the estimate error for the parameter Y_j can be calculated as,

$$\hat{E}_j = \|\hat{Y}_j - Y_j\| / \|Y_j\| \quad (4.2)$$

This method helps us to assess the change in parameter estimate over the testing/ground-truth values.

4.3 Experiments and Results

4.3.1 Data

We perform 4 sets of experiments to assess our proposed method: (1) training and testing on simulated spectra (**Synthetic - Synthetic (Spectra)**), (2) training and testing on human in-vivo spectral data from different subjects but having the same acquisition protocol (**Real (Spectra) - Real (Spectra)**), (3) training and testing on human in-vivo spectral data from different subjects with different acquisition protocol (**Real (Spectra) - Real (MRS Images)**) and (4) using the simulated spectra model to test on MRS images (**Synthetic (Spectra) - Real (MRS Images)**).

Synthetic (Spectra). A metabolite basis set was generated by using the data provided by the ISMRM MRS Fitting Challenge 2016. These were then used to simulate over 1 million spectra. In order to ensure that the simulated spectra was as close as possible to human in-vivo spectra, we incorporate the following features: variations in NAA, Cho, Cr concentrations, macro-molecular baseline, lipids, t_2 values (for changes in linewidth) and signal-to-noise ratio (SNR) to account for changes in spectral quality. As a preliminary case study, we only simulate the major metabolites (NAA, Cho and Cr) as these are easily detected by the LCModel and would, therefore, help us to evaluate the outcome of our approach and allow a suitable comparison with the LCModel. A set of over 10,000 independent test spectra were also simulated with varying combinations of the aforementioned features. For both the training and testing sets, we used the basis-set metabolite concentration values as our ground-truth.

Real (Spectra). To evaluate our method on in-vivo data, we utilize LCModel-fitted single-voxel spectroscopy (SVS) data from 287 independent human subjects. The data was obtained using the same standardized imaging protocol with the following acquisition parameters: TE/TR = 35/2000 ms, spectral width = 2500 Hz, number of points = 1024. We implement a K-fold cross-validation with 10 folds along with the random-forest regression to generate different training and testing sets having spectra from 259 and 28 subjects respectively. The metabolites assessed were: NAA, Cho, mI and Glx.

Real (MRS Images). To further assess our approach, we acquire a standard phase-encoded 2D brain MRSI data of a healthy human volunteer on a 3T scanner using a point-resolved spin-echo localization sequence (PRESS) with voxel size = 10x10x15 mm³, TE/TR=35/1000 ms, spectral width = 2000 Hz, number of points = 400. For testing purposes, we use 96 spectra from the inner-region of the brain which serves as the region of interest.

Due to the differences in acquisition parameters of the training and testing set, both the resulting spectra vary in amplitude and metabolite peak alignment. We perform a pre-processing spectral alignment step where all the test spectra are cropped from 4.3 to 0.2 ppm and interpolated to the same number of points as the training spectra to compensate for differences in acquisition bandwidth. This is followed by normalizing the amplitude of the test spectra using one of the training spectra as reference.

4.3.2 Results

Synthetic - Synthetic (Spectra). We perform an initial experiment to determine the out-of-bounds (OOB) error using different number of trees and features on a set of 20,000 simulated train and test spectra. Based on the results shown in Fig. 4.2, we proceed with the parameter estimation experiment by identifying the appropriate number of trees and features required to achieve convergence of the OOB error. For the regression error estimates, we use metabolite concentration ratios with respect to Cr (used as a standard assessment method in MRS as a means for calibration). We obtain R scores of 0.968 and 0.962 for NAA/Cr and Cho/Cr values respectively. The corresponding figures representing the linear regression are shown in Fig.4.3 and the error plots in comparison with the LCModel are shown in Fig.4.4.

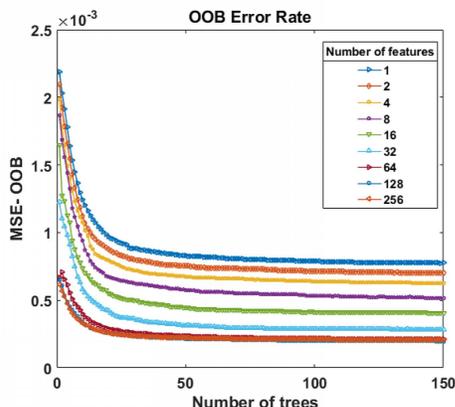


Figure 4.2: Out-Of-Bound (OOB) Error for Simulated Spectra. The experiment is performed for a varying number of features (from 1 to 256 as shown in the legend) and each iteration is assessed for a varying number of trees (as shown in the X-axis). The Y-axis represents the OOB Error rate. The error rate is minimal for more 64 features and also converges when the number of trees is close to 100.

Real (Spectra) - Real (Spectra). For the SVS dataset, we use the LCModel concentration ratio estimates as the ground-truth. Table.?? indicates the mean metabolite concentration estimate error across the 10-folds of the cross-validation process using the random forest regression method. Median error for the NAA/Cr estimate is 0.068, 0.072 for the

4.3. Experiments and Results

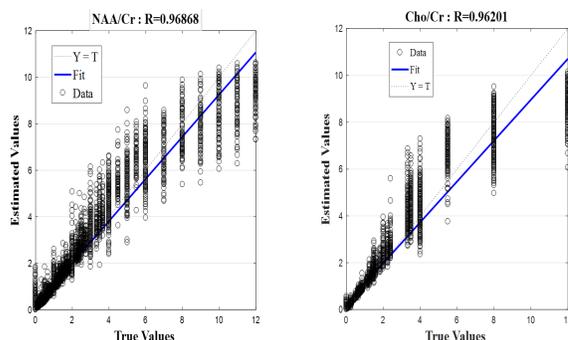


Figure 4.3: Regression Scores for the following parameters (from left to right): NAA/Cr concentration estimate and Cho/Cr concentration estimate. The X-axis represents the true values of the parameter while the y-axis represents the estimated values. Both sets of values are plotted using linear regression.

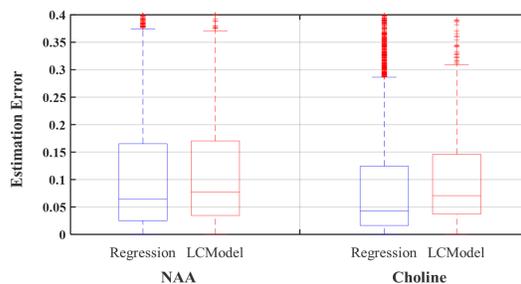


Figure 4.4: **Synthetic-Synthetic (Spec)**: Estimation error for different metabolite concentration ratios. Whiskers span the [min max] values. Median error values are represented by the red line and are as follows: NAA/Cr Regression = 0.064, NAA/Cr LCModel = 0.077, Cho/Cr Regression = 0.043, Cho/Cr LCModel = 0.070.

Cho/Cr estimate, 0.093 for the mI/Cr estimate and 0.070 for the Glx/Cr estimate compared to the corresponding LCModel estimates. The difference in error estimates is small and shows a similarity in assessment between our proposed method and the LCModel. Moreover, the low-concentration metabolites such as mI and Glx usually display a fitting error with the LCModel and the estimation error for these metabolite ratio concentrations is lower indicating that our model works well for these metabolites as well.

4. QUANTIFICATION OF METABOLITES IN MAGNETIC RESONANCE SPECTROSCOPIC IMAGING USING MACHINE LEARNING

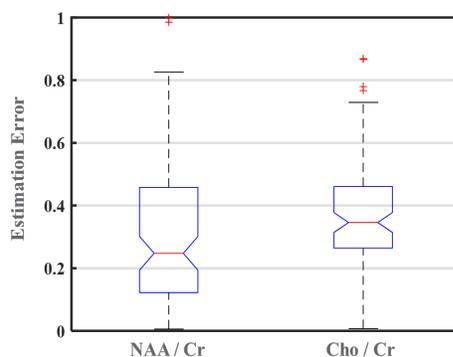


Figure 4.5: **Synthetic (Spec)-Real (MRS Images)**: Estimation error for different metabolite concentration ratios. Whiskers span the [min max] values. Median error values are represented by the red line and are as follows: NAA/Cr = 0.024, Cho/Cr = 0.034.

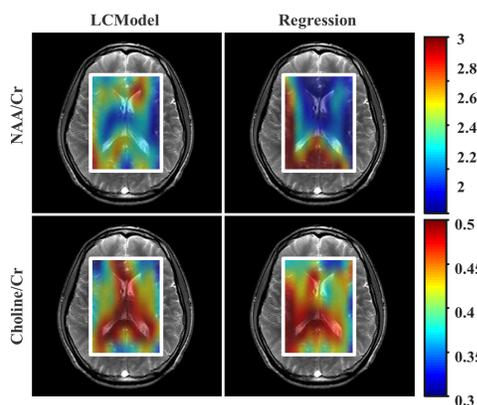


Figure 4.6: **NAA/Cr** and **Cho/Cr** concentration distribution estimates from random forest regression and non-linear model fit.

Synthetic (Spectra) - Real (Images). We test our synthetic spectra training model on the 2D MRSI data and the results are shown in the boxplot in Fig. 4.6 along with the resulting concentration distribution from both the regression approach and the non-linear model fit. As our synthetic model is trained for only NAA and Cho ratios, we show the errors for these two only. Median estimate error for NAA/Cr is 0.24 using regression. For Cho/Cr, the estimation error is 0.34. The corresponding concentration

4.3. Experiments and Results

	Naa/Cr	Cho/Cr	mI/Cr	Glx/Cr
Real-Real (Spectra)	0.068	0.072	0.093	0.070
Real-Real (Images)	0.1	0.18	0.217	0.13

Table 4.1: Concentration-ratio estimate errors using random forest regression. Results are for the experiments **Real(spec)-Real(spec)** and **Real(spec)-Real(images)**. The errors are calculated over the respective LCModel estimates as per the formula given in Eq.A.2. The major metabolites (**NAA** and **Cr**) show a low error while the smaller concentration metabolites (**mI** and **Glx**) show a slightly higher error.

Relative error	Naa	PCh
MLP	10.4%	13.2%

Table 4.2: Concentration-ratio estimate errors using random forest regression. Results are for the experiments **Real(spec)-Real(spec)** and **Real(spec)-Real(Images)**. The errors are calculated over the respective LCModel estimates as per the formula given in Eq.A.2. The major metabolites (**NAA** and **Cr**) show a low error while the smaller concentration metabolites (**mI** and **Glx**) show a slightly higher error.

values estimated from the LCModel serves as our ground-truth.

Real (Spectra) - Real (Images). We perform a blind test with 96 2D MRSI spectra against the training model generated using the 287 SVS spectra and the results are shown in Table. 4.1 and 4.2. Median estimate error for NAA/Cr is 0.1, for Cho/Cr is 0.18, for mI/Cr is 0.217 and for Glx/Cr is 0.13. Although we expect the errors to be higher in the blind test due to difference in the acquisition protocols of the training and testing dataset, the errors appear to be within a reasonable window. As expected, the estimated errors are highest for mI/Cr while Glx/Cr surprisingly has a lower error than Cho/Cr.

The Real Spectra training model provides a marginally better metabolite concentration estimate than the Synthetic spectra model . We attribute this to the presence of arbitrary scanning effects and artifacts in the real spectra model as compared to the synthetic model. For future experiments, this provides the scope for learning on a large synthetic spectral data-set with similar additional arbitrary effects to have a robust classifier for real

data (especially in the cases where annotating training data is expensive).

4.4 Conclusion

Machine learning techniques such as Random Forest-based regression provide a new and faster way of metabolite quantification. Our synthetic training model accounts for spectral features such as macro-molecular baseline, lipids, linewidth and SNR variations in combination with different metabolite concentrations. Additional features such as frequency and/or phase-shift effects along with B0 inhomogeneity could be incorporated in the model to improve robustness. For the human in-vivo data, we use training spectra from different subjects and the random-forest regression provides a low amount of estimation error over the LCModel fit even in the presence of arbitrary scanning effects. Training times for the simulated spectra can be considerable (around 5-6 hours) given that we generate over 1 million spectra while it is only a few minutes for the in-vivo spectra. On the other hand, testing and concentration estimation happens in only a few seconds and is considerably faster than the non-linear model fitting. The machine learning approach may be used directly, or indirectly by initializing LCModel fits thereby improving their results in the presence of noise and speeding up convergence. They can also be combined with global decisions about spectral quality predicting whether a spectrum can or cannot be interpreted by the physics model because of the presence of artifacts.

Future work would involve using a more robust approach such as deep-learning based methods to improve the accuracy of parameter estimation. Once a framework has been established, further work can be done on having disease-based training models for parameter estimation to predict disease progression and the corresponding metabolite maps.

Acknowledgements. The research leading to these results has received funding from the European Union’s H2020 Framework Programme (H2020-MSCA-ITN-2014) under grant agreement n° 642685 MacSeNet.

Direct Estimation of Model Parameters in MR Spectroscopic Imaging using Deep Neural Networks

This chapter has been published as **peer-reviewed conference abstract**.

© Dhritiman Das (published as open-access)

D. Das, E. Coello, A. Sekuboyina, R. F. Schulte, and B. H. Menze. “Direct estimation of model parameters in MR spectroscopic imaging using deep neural networks.” In: *Proc Intl Soc Mag Reson Med(2018), Paris, France*. 2018

Synopsis: This work deals with the metabolite quantification problem in [MR Spectroscopic Imaging \(MRSI\)](#). We introduce a multi-layer perceptron based architecture for performing regression on magnetic resonance spectra for the purpose of estimating metabolite concentration levels. Our approach is evaluated on healthy human brain MR spectra and we benchmark these results against the LCModel to evaluate improvement in speed and prediction error..

Contributions of thesis author: algorithm design and implementation, computational experiments and composition of manuscript.

5.1 Abstract

We introduce a deep-learning based framework based on a multilayer perceptron for estimation of the output parameters of a model-based analysis of MR spectroscopy data. Our proposed framework: (1) learns the spectral features from a training set comprising of different variations of synthetic spectra; (2) uses this learning and performs non-linear regression for the subsequent metabolite quantification. Experiments involve training and testing on simulated and in-vivo human brain spectra. We estimate parameters such as metabolite-concentration ratios and compare our results with that from the LCModel.

5.2 Introduction

Quantification of MR Spectroscopy (MRS) signals generates metabolic maps which show the concentration of metabolites in the sample being investigated. Accurate quantification of these metabolites is important for diagnosis of brain tumor and other in-vivo diseases. For this purpose, non-linear model-fitting tools are widely used (such as the LCModel [7], TARQUIN [10], AMARES [9] and ProFit). The LCModel is widely regarded as the gold-standard fitting tool. However, some of its drawbacks include: (1) prior knowledge-tuning and long fitting times, and (2) high estimation error for noisy data. Prior work has also focused on using machine-learning for metabolite-quantification [20]. In this study, we present an alternative to the non-linear model fitting using a deep-learning approach.

5.3 Methods

A multilayer perceptron(MLP)[26] is a fully-connected, feedforward deep-neural network comprising of three or more layers of non-linearly activated nodes. The nodes in each layer are connected to the next layer with certain weights and a supervised learning technique (backpropagation) [26] is used for training. Weights are updated after each backward-pass and the error (loss function) is computed after each iteration. Once the error reduces and achieves convergence, the learning stops.

In MRSI, The time-domain complex signal of a nucleus is given by:

$$S(t) = \int p(\omega)\exp(-i\Phi)\exp(-t/T_2^*)dw. \quad (5.1)$$

, and the corresponding frequency-domain spectrum is given by $S(\omega)$.

Using the MLP-framework, we perform the inverse signal modeling where we have a training dataset $D = (S_i(\omega), Y_i)$, $i \in [1, N]$, where N is the total number of synthetic training spectra. $S_i(\omega)$ represents the synthetic training spectral data while Y_i represents the corresponding multi-parameter training labels. As a preliminary study, for our model, we consider the concentrations (with respect to Creatine) for the major metabolites - NAA and Choline. Therefore, for a given spectrum $S_i(\omega)$, $Y_i = [NAA_i, Cho_i]$.

A five-layered perceptron network was constructed to work as a regressor mapping the $S_i(\omega)$ to the Y_i . Each layer consisted of 300 neurons with rectified linear unit (ReLU) activation. The training data consisted of $N=1$ -million spectra with their corresponding parameters. The randomly initialized network was trained to predict the parameters by iteratively minimizing the squared-error loss between the predicted and actual parameters using gradient descent with a learning rate of 1e-3. For faster convergence, Adam optimizer with a Nesterov’s momentum of 0.9 was employed. As the data is mostly well-behaved, the ‘early stopping’ convergence check was utilized on 0.1

To check the ability of our network to predict the parameters, we use two test-sets: synthetic and real CSI. The predicted concentrations are denoted by \hat{Y}_j . The corresponding LCMoDel fitted concentration labels Y_j serve as the ground-truth, $j \in [1, M]$ where M is the total number of test spectra.

Error Calculation. For our experiments, given the estimate \hat{Y}_j and the testing label Y_j , the estimate error for the parameter Y_j can be calculated as,

$$\hat{E}_j = \|\hat{Y}_j - Y_j\| / \|Y_j\| \quad (5.2)$$

Data. A metabolite basis set was generated by using the data provided by the ISMRM MRS Fitting Challenge 2016. An example has been shown in A.2. Over 1 million spectra were simulated with variations in NAA, Cho, Cr concentrations, macro-molecular baseline, lipids, t2 values (for changes in linewidth) and signal-to-noise ratio (SNR) to account for changes in spectral quality. For testing, we acquire a standard phase-encoded 2D

5. DIRECT ESTIMATION OF MODEL PARAMETERS IN MR SPECTROSCOPIC IMAGING USING DEEP NEURAL NETWORKS

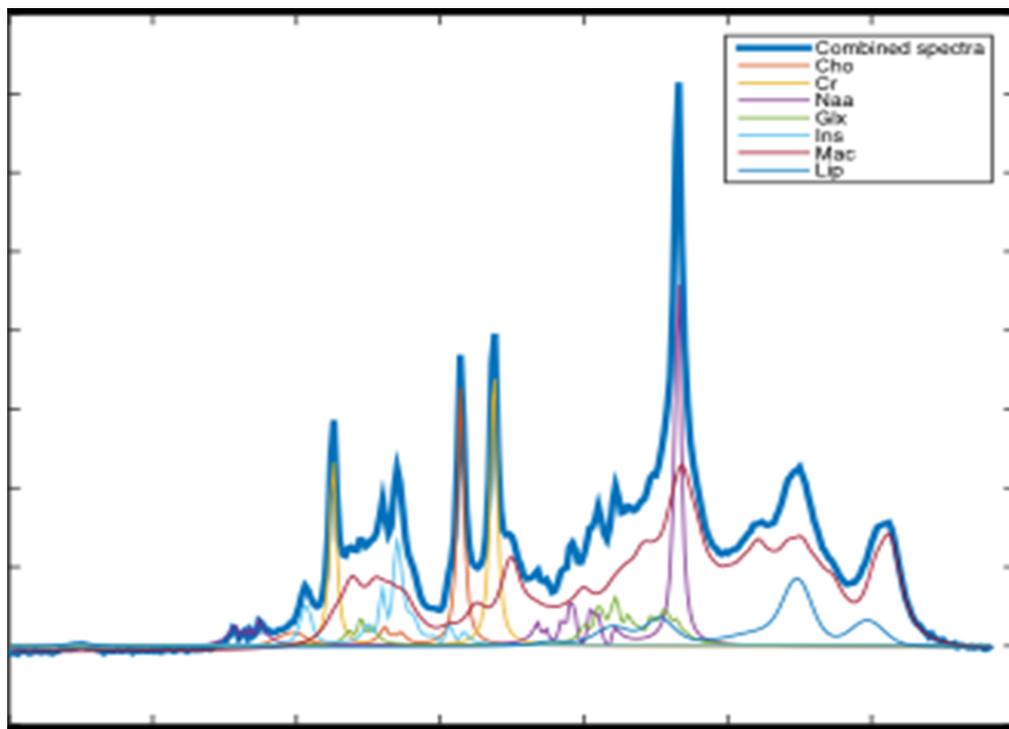


Figure 5.1: An example spectra generated using the basis sets provided by the ISMRM MRS Fitting Challenge 2016. Using the same basis sets, over **1 million** spectra are generated with variations in NAA, Cho, Cr and other metabolite concentrations along with changes in macro-molecular baseline, lipids, linewidth (t_2) and SNR.

brain MRSI data of a healthy human volunteer on a 3T scanner using a point-resolved spin-echo localization sequence (PRESS) [67] with voxel size = $10 \times 10 \times 15$ mm³, TE/TR=35/1000 ms, spectral width = 2000 Hz, number of points = 400. For testing purposes, we use 96 spectra from the inner-region of the brain which serves as the region of interest. For peak alignment, ppm-cropping and signal-normalization of the training and test spectra, a pre-processing step is performed.

5.3.1 Results

As a direct comparison of both the MLP and LCModel methods, we use the synthetic test dataset to generate the error distribution shown in Fig.

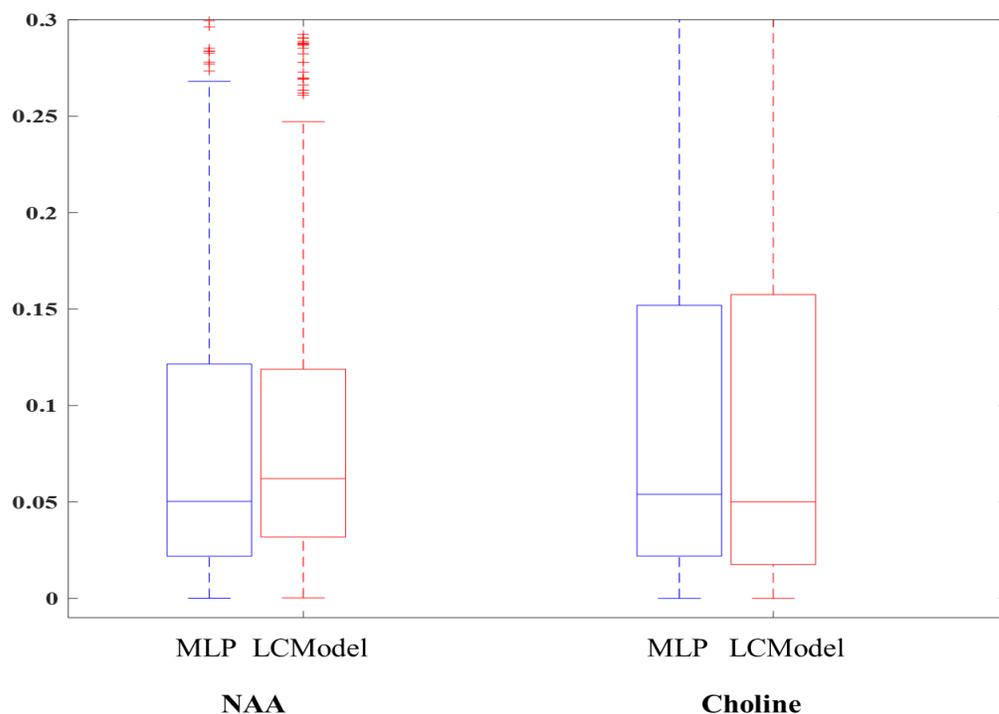


Figure 5.2: Estimation error for different metabolite concentration ratios in a synthetic spectra test-set. Whiskers span the [min max] values. Median error values are represented by the red line and are as follows: NAA/Cr MLP = 0.050, LCMoDel = 0.065, Cho/Cr MLP = 0.050, LCMoDel = 0.050.

5.2. For both NAA/Cr and Cho/Cr, the MLP shows a lower median error than the LCMoDel. Using the Bland-Altman method [68], we observe a strong correlation between the LCMoDel and RF estimates for a sample patient (Fig. A.3). \hat{E}_j for the same sample patient are within the acceptable range (especially for the major metabolites such as NAA, Choline and Creatine). Fig. 5.4 shows the resulting concentration distribution from both the MLP and LCMoDel methods for both NAA/Cr and Cho/Cr. The mean relative errors over the LCMoDel for NAA/Cr and Cho/Cr are 0.31 and 0.12 respectively.

Speed: Training time for the synthetic data is 10 minutes using the MLP. While the LCMoDel takes 10 minutes for the in-vivo metabolite

5. DIRECT ESTIMATION OF MODEL PARAMETERS IN MR SPECTROSCOPIC IMAGING USING DEEP NEURAL NETWORKS

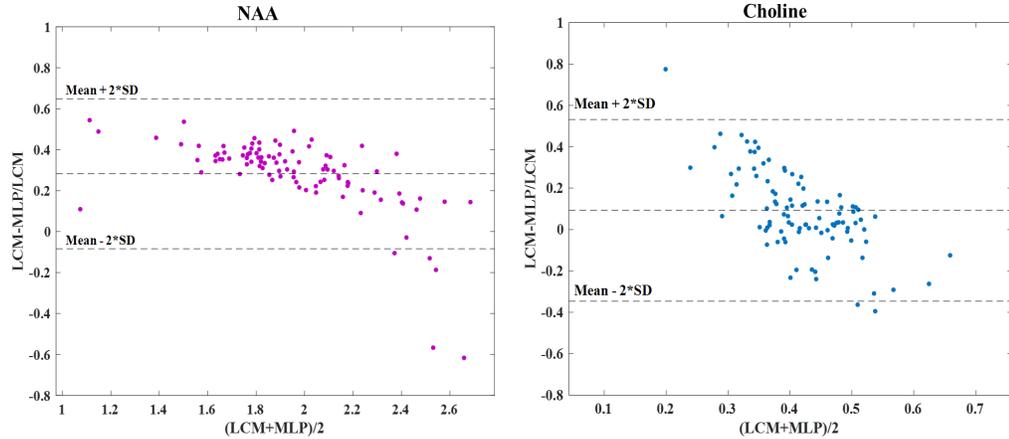


Figure 5.3: Bland-Altman plots [68] representing LCModel and Multi-Layer Perceptron (MLP) estimates of spectra for the real CSI dataset. The X-Axis is the mean of the LCModel and MLP estimate, while the Y-Axis represents the relative error of the MLP estimate over the LCModel. Bland-Altman plots for: **(Left)**: NAA and **(Right)**Choline. Both plots show a good correlation with very few outliers.

quantification, our proposed network, after training, takes only 10 seconds leading to a 60x improvement in speed.

5.4 Discussion and Conclusion

While the synthetic test-results gave a lower error compared to the LCModel, the in-vivo testing gave a slightly higher relative error. A larger training set with more training labels and a stronger network would solve this issue by providing a robust classification of real data. In our proposed method, testing and concentration estimation happens in only a few seconds and is considerably faster than the LCModel fitting. The deep neural-networks may be used directly, or indirectly by initializing LCModel fits thereby improving their results in the presence of noise and speeding up convergence.

Future work would involve using a more diverse network with layer-wise training of spectral features to improve the accuracy of parameter estimation. Once a framework has been established, further work can be done on combining these networks with global decisions about predicting

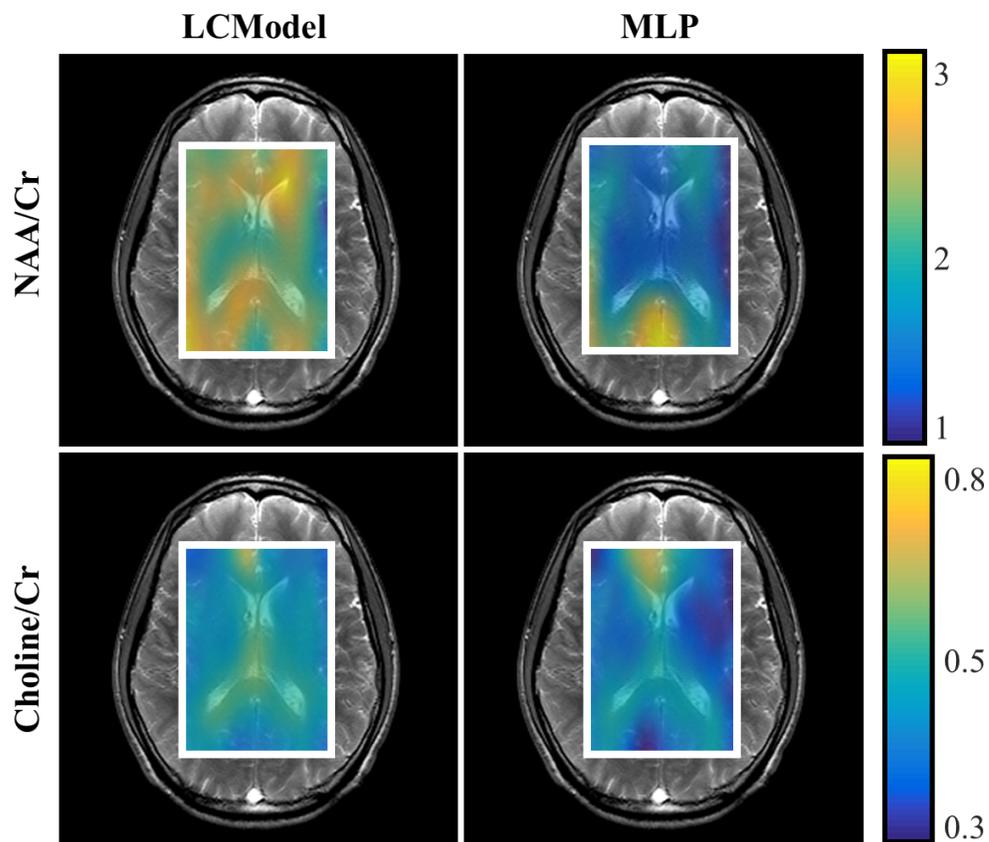


Figure 5.4: NAA/Cr and Cho/Cr concentration distribution estimates from **(Left)**LCMoDel fit and **(Right)** multi-layer perceptron (MLP). The mean relative errors are 0.31 for NAA/Cr and 0.12 for Cho/Cr.

spectral quality especially in the presence of artifacts.

Acknowledgements. The research leading to these results has received funding from the European Union’s H2020 Framework Programme (H2020-MSCA-ITN-2014) under grant agreement n° 642685 MacSeNet.

Concluding Remarks

[MR Spectroscopic Imaging \(MRSI\)](#) is a clinical imaging modality for studying tissues in-vivo to investigate and diagnose neurological diseases. More specifically, this modality can be used for measuring specific tissue metabolites and thereby enable the non-invasive diagnosis and characterization of patho-physiological changes in tissues, especially the brain, and help in early detection of highly proliferative disease such as glioblastoma. While the modality has been around for a long time dating back to the discovery of [nuclear magnetic resonance \(NMR\)](#) imaging, its implementation clinically has been severely limited due to several technical challenges. Over the course of the last few decades it has been overtaken by proton MR imaging, diffusion MRI and other imaging methodologies which facilitate ease-of-use in a research and clinical environment.

Recent progress of advanced biomedical image analysis techniques facilitated by state-of-the-art computational methods offer an avenue for dealing with these challenges and leveraging the diagnostic power of [MRSI](#). In this dissertation, we sought to provide computational solutions to two of the major challenges while aiming to accelerate the analysis and processing workflow for the MRS data by proposing advanced *data-driven* solutions. Our contributions focus on tackling two prominent challenges - denoising and metabolite quantification. Due to the publication-based nature of this thesis, the [Chapters 3 to 5](#) are self-contained and in their original form. [Appendix A](#) and [Appendix B](#) provide additional, unpublished work which may evolve in the future. This final chapter, therefore, provides an overview of the preceding chapters as well as a general discussion of the application and finally concluding with the directions for future research.

First, we have presented a novel denoising method specifically for MRSI data in [Chapter 3](#). This method addresses the issue of arbitrary phase,

6. CONCLUDING REMARKS

induced in the data during [MRSI](#) acquisition, by generating a frequency-phase map for each spectra in the data. This involves a *dephasing* step to incorporate a set of phase-shifts in the acquired spectra through a pre-defined range thereby artificially inducing redundancy in the data even in the presence of arbitrary, non-parametric phase thereby aiding in the patch-based denoising method. There are a few remarks on this part of our work. First, as with any denoising method based on patch-based averaging, it is important to ensure that there is no over-smoothing induced in the data which may lead to loss of important metabolite information. While we were able to successfully demonstrate that our method retains the spatial-spectral resolution while improving the SNR, this was evaluated on only the major metabolites. Future work would also involve incorporating all the available metabolites as part of the clinical pipeline (including those which have overlapping resonance peaks) and thereby improve the signal quality. Our method was also evaluated on short-TE spectra acquired from a 3T-scanner. A more interesting direction may involve validating this approach on higher resolution, spatially-resolved spectra acquired from scanners having a higher magnetic field, such as 7T or 9.4T, which could overcome the issues of overlapping resonance peaks and presence of artifacts.

In Chapter 4, we have proposed a machine-learning based regression approach using random forests for metabolite quantification. The goal of developing this architecture was to provide an alternative to the conventional, model fitting tools which were being used by the MRS community. The key challenges to address involved the long fitting times, the need for manual parameter tuning and the high estimation error rates in the presence of noisy artifacts and peak broadening. In hindsight, this wound up being one of the first attempts to use a machine-learning approach for quantification of metabolites in MR spectra and eventually paved the way for rapid development in this aspect of [MRSI](#) processing and analysis. One of the key innovations in this work involved generating a training set of over 1 million synthetic spectra with variations in metabolite concentrations, frequency-phase shifts, noise levels and macromolecular baseline artifacts. The fact that we used the individual metabolite basis sets in our simulation model (based on [NMR](#) physics) allowed the generation of synthetic spectra which was somewhat representative of human, in-vivo brain spectra. Our proposed architecture not only gave comparable performance with the state-of-the-art fitting tool, the LCModel, but also performed the quantification in a fraction of the time. As a comparison, for a 32x32 healthy human

brain MRS dataset with 256 spectral points, the LCModel took almost 1 hour to perform the fitting whereas our approach took about 10 seconds after training while giving similar concentration estimates. In Appendix A and Appendix B, we performed additional experiments on clinical datasets involving patients with secondary-progressive multiple sclerosis and glioma respectively, while also exploring other architectures such as the [multi-layer perceptron \(MLP\)](#) and [convolutional neural network \(CNN\)](#). The potential for this approach is exciting, however, a lot more work still needs to be done. We have evaluated our method on only the major macro-metabolites which are used for assessment and diagnosis of neurological diseases. For this work to transition to a clinically usable tool, it would have to involve quantification of all the metabolites (including the smaller ones) in a conventional [MRSI](#) protocol. A visual representation of the metabolite signals post-quantification along with the baseline signals would aid the clinicians in interpreting the results of this tool.

As with many other fields, recent progress in machine learning has strongly influenced the latest methods being developed in medical image analysis. A quick look at the majority of methods used across popular medical imaging challenges, such as the brain tumor segmentation challenge (BRATS) [69] or the fastMRI challenge [70] employ deep-learning based methods as part of their workflow. This has led to a drastic improvement in the reconstruction and analysis of complex, medical imaging data especially in the presence of artifacts or other abnormalities. Given the rapid pace of progress in machine-learning theory and its application, one can expect more developments in the coming years: adapting the next generation of machine learning and computer vision methods for metabolite quantification and spectral quality assessment especially for clinical conditions such as glioblastoma, multiple sclerosis, among others, will most likely remain a main direction of research that will only serve to benefit the clinical application of [MRSI](#).

Since the formulation of this thesis, there have been a few projects done across multiple-sites which have explored [CNN](#)-based architectures to integrate both the denoising of spectra and metabolite quantification. The use of Generative Adversarial Networks (GAN) [26] for generating realistic synthetic spectra, if successful, can drastically improve the training process and lead to improved concentration prediction outcomes. Bayesian inference and federated-learning methods are also being used for other applications and, if applied to [MRSI](#), can serve to improve the overall spectral acquisition

6. CONCLUDING REMARKS

and analysis pipeline. The eventual goal would be to develop an open-source, easily accessible analysis suite for [MRSI](#) which would enable clinicians and researchers to make better use of this modality for disease assessment and diagnosis.

In conclusion, with this work, we have addressed some of the major processing challenges that have hindered the widespread clinical use of [MRSI](#). With continued improvements aided by developments in the field of computational approaches and more efficient data-driven methods, one hopes that [MRSI](#) may become a part of the standard clinical workflow and be integrated with other imaging methodologies to improve patient outcomes.

Appendices

Metabolite Quantification of ^1H MRSI spectra in Multiple Sclerosis: A Machine Learning Approach

Contributions of thesis author: Unpublished work presented as part of the thesis.

A.1 Abstract

As an alternative to model-based spectral fitting tools, we introduce a machine-learning framework for estimating metabolite concentrations in MR spectra acquired from a homogeneous cohort of 42 patients with Secondary Progressive Multiple Sclerosis. Our framework based on random-forest regression performs a 42-fold cross validation on this dataset which involves (1) learning the spectral features from this cohort; (2) estimating concentrations and calculating relative error over the LCModel estimates. Compared to the LCModel, our method, after training, gives a low estimation error and a 60-fold improvement in estimation speed per patient.

A.2 Purpose

Multiple Sclerosis (MS) is an inflammatory neurological diseases affecting around 2.5 million people globally (www.mstrust.org.uk). Quantification of metabolites can yield important diagnostic information about disease

A. METABOLITE QUANTIFICATION OF ^1H MRSI SPECTRA IN MULTIPLE SCLEROSIS: A MACHINE LEARNING APPROACH

progression for which model-based fitting tools such as the LCModel [7] are widely used. Despite being the gold-standard, they require prior knowledge tuning and can also be time-consuming. Prior work on healthy human brains has shown Random-Forest regression to perform the same quantification but at a faster speed [20]. The purpose of this work is to apply this method to a clinical setup involving MS patients and address the limitations of parameter tuning and speed of conventional fitting tools by validating on an extensive ^1H MRSI-MS dataset.

A.3 Methods

Random Forests [66] have been shown to be effective in a wide range of classification [60][59] and regression problems [20], including classification of MS spectra [71]. These involve multiple forests comprising of a set of binary trees. For training, splits are created in each tree based on random subsets of the feature variables and piecewise linear regression is performed over the input data [20]. The process involves seeking best prediction at every node and using thresholding to further propagate data points till they reach the end of the tree. Subsequently the weighted average of the prediction from each tree is taken to give a single output estimate.

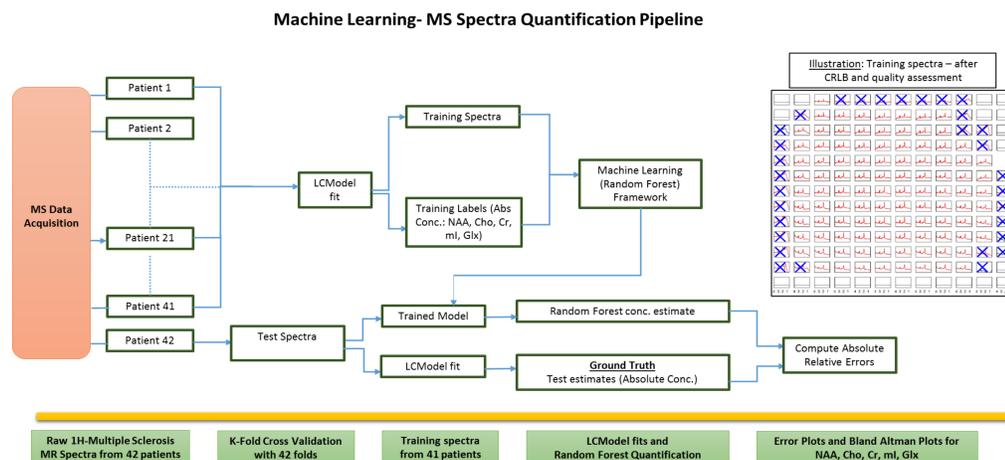


Figure A.1: Pipeline for Machine Learning based quantification of MS spectra. Using the K-fold cross validation approach, this method is repeated using each of the 42 patient MS-spectra as test spectra.

In MRSI, The time-domain complex signal of a nucleus is given by:

$$S(t) = \int p(\omega) \exp(-i\Phi) \exp(-t/T_2^*) d\omega. \quad (\text{A.1})$$

, and the corresponding frequency-domain spectrum is given by $S(\omega)$.

As shown in Fig. A.1, we aim to perform the inverse signal modeling where we have a training dataset $D = (S_i(\omega), Y_i)$, $i \in [1, N]$, where N is the total number of MS training spectra from 41 patients. $S_i(\omega)$ represents the training spectral data while Y_i represents the corresponding multi-parameter training labels. For our model, we consider the absolute concentrations of NAA, Cho, Cr, Myo-Inositol (mI) and Glutamate+Glutamine (Glx). Therefore, for a given spectrum $S_i(\omega)$, $Y_i = [\text{NAA}_i, \text{Cho}_i, \text{Cr}_i, \text{mI}_i, \text{Glx}_i]$.

We run a K-Fold cross validation with 42 folds, number of trees = 100 and mTry = 128 to generate training sets using spectra from 41 patients and test on the spectra, $S_j(\omega)$, from the remaining patient to obtain concentration estimates \hat{Y}_j . The corresponding LCModel fitted concentration labels Y_j serve as the ground-truth, $j \in [1, M]$ where M is the total number of test spectra.

Error Calculation. For our experiments, given the estimate \hat{Y}_j and the testing label Y_j , the estimate error for the parameter Y_j can be calculated as,

$$\hat{E}_j = \|\hat{Y}_j - Y_j\| / \|Y_j\| \quad (\text{A.2})$$

This method helps us to assess the absolute relative change in parameter estimates over the ground-truth values.

Subjects 42 patients (age 55+/-8) with SPMS (EDSS score 6+/-0.7) were recruited as part of the study. The patients underwent MRI examination at 3T, including T1w, T2w and FLAIR scans and proton semi-LASER MRSI with TR/TE of 2000/43 ms and 481 data points over spectral region-of-interest. An example has been shown in Fig. A.2. Prior work provides additional acquisition details [71]. The data was fitted using LCModel and the voxels which did not pass the Cramer-Rao Lower Bound (CRLB) and visual quality tests were discarded (e.g. distorted baselines). There are no 'pure' lesion or non-lesion voxels and therefore all spectra are of mixed content. Total spectra per fold included approximately 4200 training and 100 test-spectra/patient respectively.

A. METABOLITE QUANTIFICATION OF ^1H MRSI SPECTRA IN MULTIPLE SCLEROSIS: A MACHINE LEARNING APPROACH

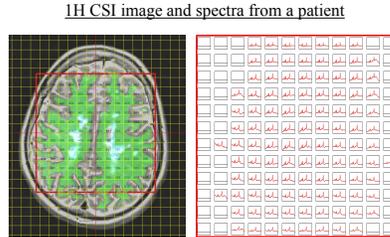


Figure A.2: MRSI data acquired from an SPMS patient and the corresponding spectra after voxel-wise LCModel fitting [71].

A.4 Results

We aim to get an error \hat{E}_j of $< 15\%$ for each of the metabolites per patient. The mean absolute relative error plots for all patients are shown in Fig. A.4. Using the Bland-Altman method [68], we observe a strong correlation between the LCModel and RF estimates for a sample patient (Fig. A.3). \hat{E}_j for the same sample patient are within the acceptable range (especially for the major metabolites such as NAA, Choline and Creatine) (Fig. A.5). **Speed:** Training time per fold is 3 minutes. While the LCModel takes 10 minutes per patient for the metabolite quantification, our proposed framework, after training, takes only 10 seconds leading to a 60x improvement in speed.

A.5 Discussion

An outlier high-error spectrum for a test patient has been highlighted in Fig. A.3 which corresponds to the outlier whisker in Fig. A.5. The peaks in this spectrum are not well-defined and only few such spectra are present in the training set. The machine learning algorithm is therefore insufficiently trained on such spectra resulting in a high error.

Future work would involve using a more robust approach such as deep-learning based methods to improve the accuracy of parameter estimation. Once a framework has been established, further work can be done on

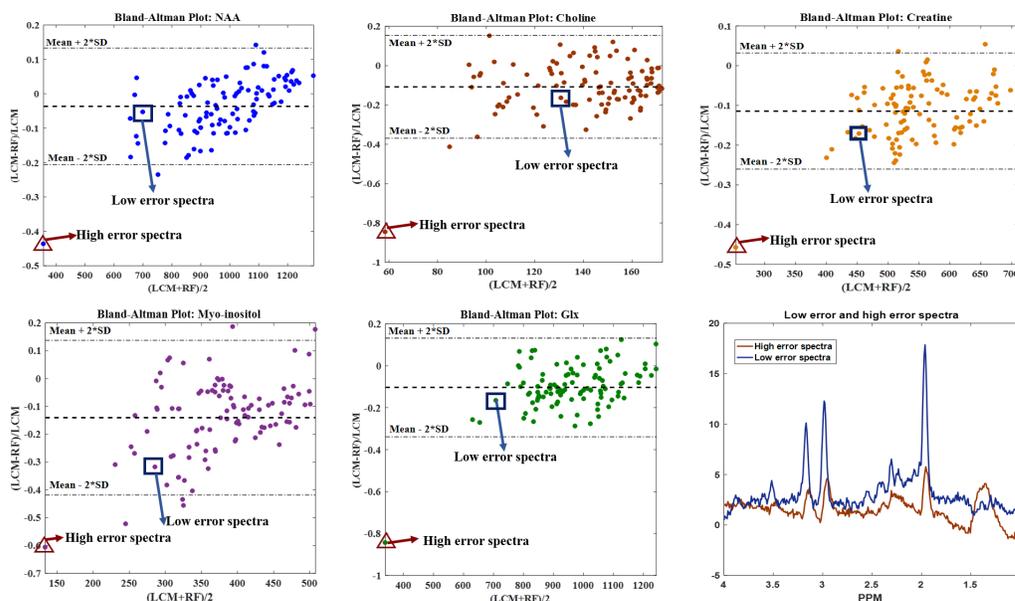


Figure A.3: Bland-Altman plots [68] representing LCModel and Random Forest (RF) estimates of spectra from a sample MS patient dataset. The X-Axis is the mean of the LCModel and RF estimate, while the Y-Axis represents the relative error of the RF estimate over the LCModel. **From (L-R):** Bland-Altman Plots for (Top row)- NAA, Choline and Creatine, (Bottom row)- myo-Inositol and Glx and spectra plots of outlier spectrum with high quantification error and one of the spectra with low quantification error (these are marked in the Bland-Altman plots within a red triangle and blue square respectively).

having a MS-specific simulated spectra model with wider variations in the spectral peak for parameter estimation to predict disease progression and the corresponding metabolite maps.

Acknowledgements. Spectral data were provided by the MS-SMART consortium (www.ms-smart.org). The research leading to these results has received funding from the European Union’s H2020 Framework Programme (H2020-MSCA-ITN-2014) under grant agreement n° 642685 MacSeNet.

A. METABOLITE QUANTIFICATION OF ^1H MRSI SPECTRA IN MULTIPLE SCLEROSIS: A MACHINE LEARNING APPROACH

Relative Mean Absolute Error for 42 patients

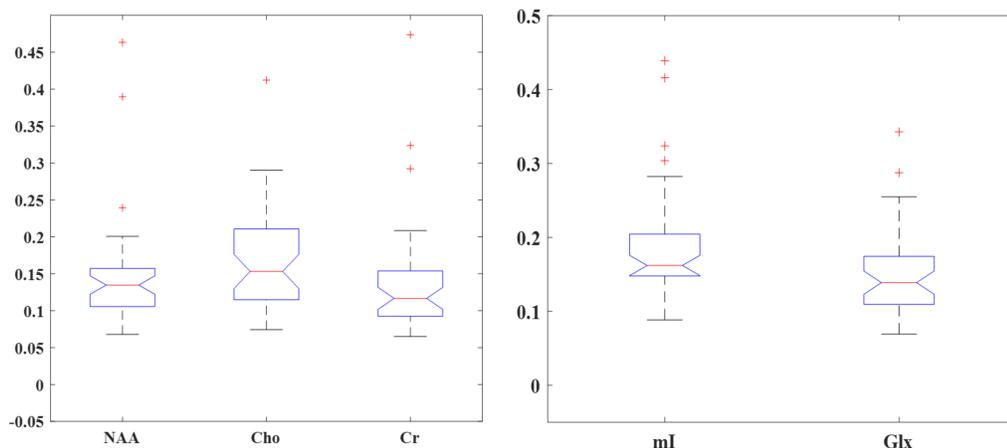


Figure A.4: Absolute Relative Error box-plots for the entire 42-patient spectra dataset. The mean relative errors (red line) are: **Left** NAA (5.4%), Choline (10.4%), Creatine (10.5%), **Right** mI (15.7%), Glx (13.8%). 8 patients had an average error of $> 15\%$ per metabolite which contributed to the slight increase in the mean error across all the patients.

Relative Mean Absolute Error for a Sample patient

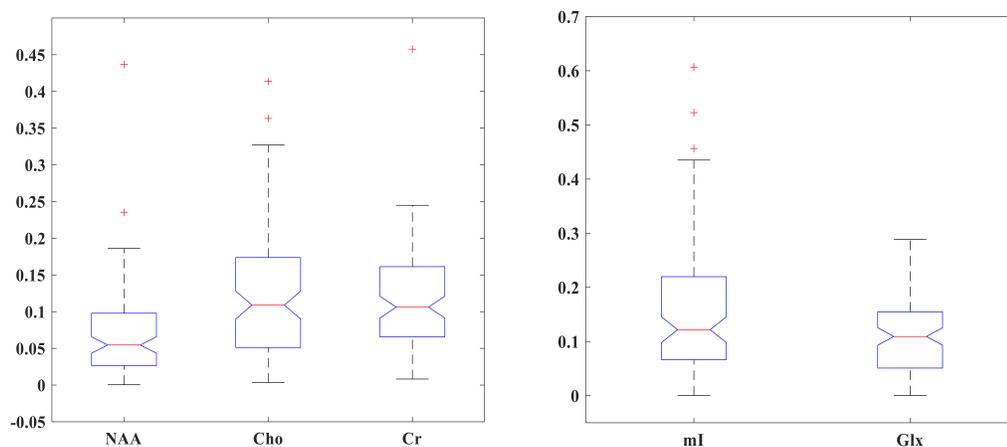


Figure A.5: Absolute Relative Error box-plots for the sample patient spectra. The mean relative errors (red line) are: (Left) NAA (13.4%), Choline (14.1%), Creatine (11.5%), (Right) mI (11.5%), Glx (10.8%).

Using Deep Neural Networks for Metabolite Quantification in Healthy and Glioma spectra

Contributions of thesis author: Unpublished work presented as part of the thesis.

B.1 Abstract

Following the work presented in 5, we developed an architecture based on 1D-convolutional neural network (CNN) for the purpose of metabolite quantification using regression. The architecture was optimized for different layers and kernel sizes and the quantification pipeline was validated on healthy, human 2D MR Spectroscopic Imaging (MRSI) dataset and also on 40 patients with different stages of *glioma*. Results were compared with previous methods including random forest (RF) and multi-layer perceptron (MLP).

B.2 Architecture design

The multi-layer perceptron (MLP) architecture comprised of 4 layers with 400 – 200 – 100 – 64 neurons. The input spectra length was 256 points and the hypothesis was to have an initial layer with the number of neurons being double than the number of spectral points in order to extract minor features from the spectra. For the purpose of training, we used an Adam

B. USING DEEP NEURAL NETWORKS FOR METABOLITE QUANTIFICATION IN HEALTHY AND GLIOMA SPECTRA

Optimizer with an initial learning rate of 0.0001. Adding an L_2 regularizer of 0.01 helped in improving our results without it.

The [convolutional neural network \(CNN\)](#) architecture was designed with 3 hidden layers (1d kernel spread=3, stride =2) and output channels of 8, 16 and 32. The final layer was an [fully convolutional neural network \(FCN\)](#) comprising of 128 neurons. For training, we, again, used an Adam Optimizer along with a leaky ReLU with static decay; average pooling of 2x2 was performed after each layer and dropout was also introduced at the end of the third layer.

B.3 Data and Experiments

For the healthy human dataset, we trained on the previously defined synthetic dataset comprising of *1 million* short-TE training spectra and evaluated on the healthy, human 2D CSI brain dataset comprising of 96 spectra. For this study, we compared performance between the [random forest \(RF\)](#), [multi-layer perceptron \(MLP\)](#) and our [convolutional neural network \(CNN\)](#) architecture.

The *glioma* dataset comprised of *long-TE* brain spectra from 40 multi-stage glioma patients. The data comprised of approximately 400 fitted spectra per patient with each spectra having 1024 points. Fig. B.1 shows an example of glioma spectra from a sample patient. We performed a patient-wise K-fold cross-validation using our [convolutional neural network \(CNN\)](#) architecture and evaluated for absolute concentrations for the major metabolites (which serve as indicators of glioma progression).

B.4 Results and Discussion

Results for the comparison study for the healthy subject are shown in Fig. B.2. Testing time took approximately 10 seconds for the entire test dataset for each of the architectures. While the [multi-layer perceptron \(MLP\)](#) performs the best, the [convolutional neural network \(CNN\)](#) gives a comparable performance though underestimating the metabolite concentration distribution for NAA. One of the key takeaways of this study is that for a synthetic training set [multi-layer perceptron \(MLP\)](#) and [random forest \(RF\)](#) seem to

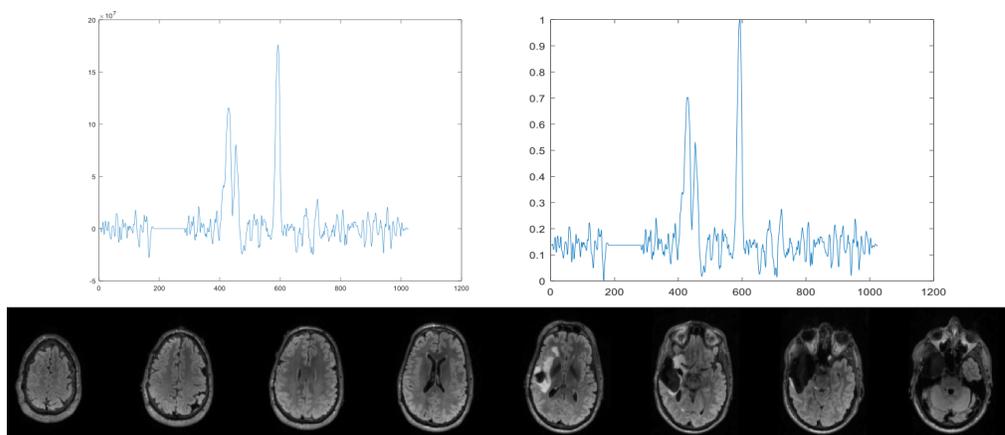


Figure B.1: Example of long-TE glioma dataset from a sample patient.

generalize better to real data as compared to [convolutional neural network \(CNN\)](#).

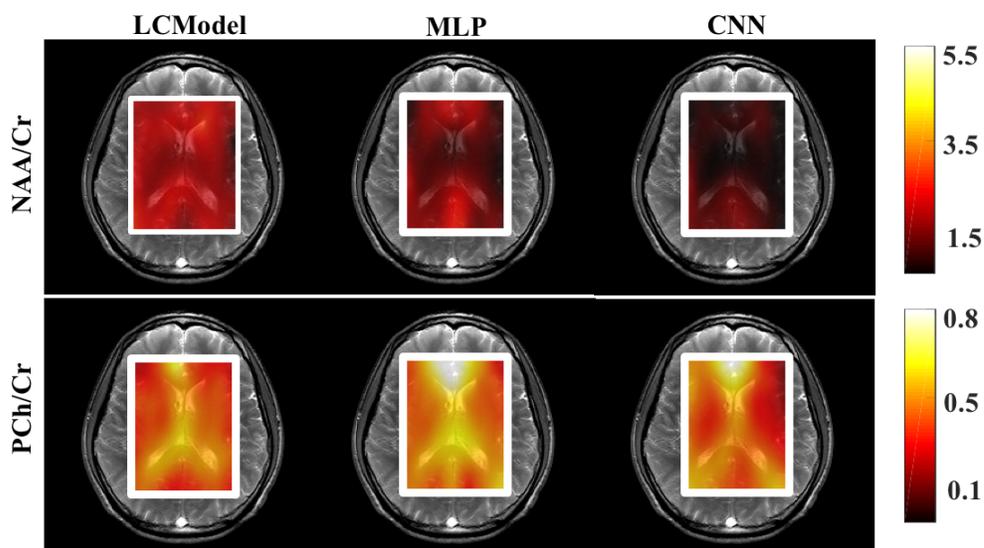


Figure B.2: Metabolite maps following quantification using 3 different architectures. The MLP gives the better distribution out of all 3 while the CNN comes close though it ends up underestimating the metabolite concentration distribution for NAA.

B. USING DEEP NEURAL NETWORKS FOR METABOLITE QUANTIFICATION IN HEALTHY AND GLIOMA SPECTRA

Architecture	Choline	Creatine	Naa
Random Forest	13.7%	11.5%	14%
MLP	11.5%	9.2%	12%
CNN	10.1%	10.8%	12.3%

Table B.1: Concentration estimate errors using random forests, multi-layer perceptrons and convolutional neural networks on 40 glioma patients. A patient-wise K-fold cross validation was performed and the results shown are averaged over all the folds. Overall, the cnn ends up out-performing the other 2 architectures in terms of metabolite quantification.

This work can be further improved by using generative models (GANs) which can help in generating realistic, synthetic dataset while also having a broader training database comprising of short and long-TE healthy and glioma spectra.

Acknowledgements. Spectral data were provided by Jason Crane and Yan Li from the University of California-San Francisco (UCSF). The research leading to these results has received funding from the European Union’s H2020 Framework Programme (H2020-MSCA-ITN-2014) under grant agreement n° 642685 MacSeNet.



List of Publications

The following publications were written *during this thesis*.

Peer-reviewed Conference Proceedings

- D. Das, E. Coello, R. F. Schulte, and B. H. Menze. “Quantification of Metabolites in Magnetic Resonance Spectroscopic Imaging Using Machine Learning.” In: *Medical Image Computing and Computer Assisted Intervention, MICCAI 2017*. Ed. by M. Descoteaux, L. Maier-Hein, A. Franz, P. Jannin, D. L. Collins, and S. Duchesne. Cham: Springer International Publishing, 2017, pp. 462–470. DOI: [10.1007/978-3-319-66179-7_53](https://doi.org/10.1007/978-3-319-66179-7_53).
- D. Das, E. Coello, R. F. Schulte, and B. H. Menze. “Spatially Adaptive Spectral Denoising for MR Spectroscopic Imaging using Frequency-Phase Non-Local Means.” In: *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2016*. Ed. by S. Ourselin, L. Joskowicz, M. R. Sabuncu, G. Unal, and W. Wells. Cham: Springer International Publishing, 2016, pp. 596–604. DOI: [10.1007/978-3-319-46726-9_69](https://doi.org/10.1007/978-3-319-46726-9_69).

Journal Articles

- C. Ulas, D. Das, M. J. Thrippleton, M. d. C. Valdes Hernandez, P. A. Armitage, S. D. Makin, J. M. Wardlaw, and B. H. Menze. “Convolutional Neural Networks for Direct Inference of Pharmacokinetic Parameters: Application to Stroke Dynamic Contrast-Enhanced MRI.”

C. LIST OF PUBLICATIONS

In: *Frontiers in Neurology* 9 (2019), p. 1147. DOI: [10.3389/fneur.2018.01147](https://doi.org/10.3389/fneur.2018.01147).

- B. Kanberoglu, D. Das, P. Nair, P. Turaga, and D. Frakes. “An Optical Flow-Based Approach for Minimally Divergent Velocimetry Data Interpolation.” In: *International Journal of Biomedical Imaging* 2019 (2019). DOI: [10.1155/2019/9435163](https://doi.org/10.1155/2019/9435163).

Peer-reviewed Conference Abstract Proceedings

- D. Das, E. Coello, R. F. Schulte, and B. H. Menze. “Spectral-Dephasing based Non-Local Means for Spatial Adaptive Denoising in 3D-MRSI.” In: *Proc Intl Soc Mag Reson Med(2017), Honolulu, Hawaii, USA*. 2017.
- D. Das, M. E. Davies, J. Chataway, S. Chandran, B. H. Menze, and I. Marshall. “Metabolite Quantification of 1H-MRSI spectra in Multiple Sclerosis: A Machine Learning Approach.” In: *Proc Intl Soc Mag Reson Med(2018), Paris, France*. 2018.
- D. Das, E. Coello, A. Sekuboyina, R. F. Schulte, and B. H. Menze. “Direct estimation of model parameters in MR spectroscopic imaging using deep neural networks.” In: *Proc Intl Soc Mag Reson Med(2018), Paris, France*. 2018.
- D. Das, M. J. Thrippleton, M. E. Davies, R. F. Schulte, S. Chandran, B. H. Menze, and I. Marshall. “MRSI Brain Temperature Mapping using Machine Learning.” In: *Proc Intl Soc Mag Reson Med(2018), Paris, France*. 2018.

D

Original Publications

The following peer-reviewed publications, in this section, were presented *as a part of this thesis* and are included here in their original published format.

Spatially Adaptive Spectral Denoising for MR Spectroscopic Imaging using Frequency-Phase Non-local Means

Dhritiman Das^{1,3}(), Eduardo Coello^{2,3}, Rolf F. Schulte³,
and Bjoern H. Menze¹

¹ Department of Computer Science, Technical University of Munich,
Munich, Germany

`dhritiman.das@tum.de`

² Department of Physics, Technical University of Munich, Munich, Germany

³ GE Global Research, Munich, Germany

Abstract. Magnetic resonance spectroscopic imaging (MRSI) is an imaging modality used for generating metabolic maps of the tissue in-vivo. These maps show the concentration of metabolites in the sample being investigated and their accurate quantification is important to diagnose diseases. However, the major roadblocks in accurate metabolite quantification are: low spatial resolution, long scanning times, poor signal-to-noise ratio (SNR) and the subsequent noise-sensitive non-linear model fitting. In this work, we propose a frequency-phase spectral denoising method based on the concept of non-local means (NLM) that improves the robustness of data analysis and scanning times while potentially increasing spatial resolution. We evaluate our method on simulated data sets as well as on human in-vivo MRSI data. Our denoising method improves the SNR while maintaining the spatial resolution of the spectra.

1 Introduction

Magnetic Resonance Spectroscopic imaging (MRSI), also known as chemical shift imaging, is a clinical imaging modality for studying tissues in-vivo to investigate and diagnose neurological diseases. More specifically, this modality can be used in non-invasive diagnosis and characterization of patho-physiological changes by measuring specific tissue metabolites in the brain. Accurate metabolite quantification is a crucial requirement for effectively using MRSI for diagnostic purposes. However, a major challenge with MRSI is the long scanning time required to obtain spatially resolved spectra due to abundance of metabolites that have a concentration which is approximately 10,000 times smaller than water. Current acquisition techniques such as Parallel Imaging [13] and Echo-Planar Spectroscopic Imaging [9] focus on accelerated scanning times combined with reconstruction techniques to improve the SNR of the spectral signal. Despite this, further accelerated acquisitions are desirable. Furthermore, an improved SNR is needed as the non-linear voxel-wise fitting to noisy data leads to a high amount of local minima and noise amplification resulting in poor spatial resolution [7].

The SNR of the signal can be improved by post-processing methods such as denoising algorithms [8], apodization (Gaussian/Lorentzian), filter-based smoothing and transform-based methods [3]. However, these methods reduce resolution and remove important quantifiable information by averaging out the lower-concentration metabolites. Recently, data-dependent approaches such as the Non-local Means (NLM), which use the redundancy inherent in periodic images, are being used extensively for denoising [3]. In the case of MRS, this periodicity implies that the spectra in any voxel may have similar spectra in other voxels in the frequency-phase space. Therefore, it carries out a weighted average of the voxels in this space, depending on the similarity of the spectral information of their neighborhoods to the neighborhood of the voxel to be denoised.

Our Contribution. *In this work, we propose a method for spectrally adaptive denoising of MRSI spectra in the frequency-phase space based on the concept of Non-local Means.* Our method compensates for the lack of phase-information in the acquired spectra by implementing a dephasing approach on the spectral data. In the next section, we introduce the experimental methods beginning with the concept of NLM in the frequency-phase space followed by the spectral dephasing and rephasing approach. Our proposed method is then validated quantitatively and qualitatively using simulated brain data and human in-vivo MRSI data sets to show the improvements in SNR and spatial-spectral resolution of MRSI data.

2 Methods

MR Spectroscopy. Magnetic resonance spectroscopy is based on the concept of nuclear magnetic resonance (NMR). It exploits the resonance frequency of a molecule, which depends on its chemical structure, to obtain information about the concentration of a particular metabolite [12]. The time-domain complex signal of a nuclei is given by: $S(t) = \int p(\omega) \exp(-i\Phi) \exp(-t/T_2^*) d\omega$. The frequency-domain signal is given by $S(\omega)$, T_2^* is the magnetization decay in the transverse plane due to magnetic field inhomogeneity and $p(\omega)$ comprises of Lorentzian absorption and dispersion line-shapes function having the spectroscopic information about the sample. Φ represents the phase, $(\omega t + \omega_0)$, of the acquired signal where ωt is the time-varying phase change and ω_0 is the initial phase. For the acquired MRSI data, I , Φ is unknown. This process allows generation of metabolic maps through non-linear fitting to estimate concentration of metabolites such as N-acetyl-aspartate (NAA), Creatine (Cr) and Choline (Cho).

2.1 Non-local Means (NLM) in Frequency-Phase Space

As proposed by Buades et al. [3], the Non-local Means (NLM) method restores the intensity of voxel x_{ij} by computing a similarity-based weighted average of all the voxels in a given image. In the following, we adapt NLM to the MRSI data: let us suppose that we have complex data, $I : \Omega^3 \mapsto \mathbb{C}$ of size $M \times N$ and noisy spectra $S_{ij}(\omega)$, where $(x_{ij} | i \in [1, M], j \in [1, N])$ and Ω^3 is the frequency-phase

grid. Using NLM for denoising, the restored spectra $\hat{S}_{ij}(\omega)$ is computed as the weighted average of all other spectra in the frequency-phase space defined as:

$$\hat{S}_{ij}(\omega) = \sum_{x_{kl} \in \Omega^3} w(x_{ij}, x_{kl}) S_{kl}(\omega) \quad (1)$$

As a probabilistic interpretation, spectral data $S_{11}(\omega), \dots, S_{MN}(\omega)$ of voxels x_{11}, \dots, x_{MN} respectively are considered as MN random variables X_{ij} and the weighted average estimate $\hat{S}_{ij}(\omega)$ is the maximum likelihood estimate of $S_{ij}(\omega)$.

$N_{ij} = (2p + 1)^3$, $p \in \mathbb{N}$ is the cubic neighborhood of voxel x_{ij} within the search volume $V_{ij} = (2R + 1)^3$ around x_{ij} along the frequency, phase and spatial directions. $R \in \mathbb{N}$, where R is the radius of search centered at the voxel x_{ij} . The weight $w(x_{ij}, x_{kl})$ serves as a quantifiable similarity metric between the neighborhoods N_{ij} and N_{kl} of the voxels x_{ij} and x_{kl} provided $w(x_{ij}, x_{kl}) \in [0, 1]$ and $\sum w(x_{ij}, x_{kl}) = 1$. The Gaussian-weighted Euclidean distance is computed between $S(N_{ij})$ and $S(N_{kl})$ as shown below:

$$w(x_{ij}, x_{kl}) = \frac{1}{Z_{ij}} e^{-\frac{\|S(N_{ij}) - S(N_{kl})\|_2^2}{h^2}} \quad (2)$$

where Z_{ij} serves as the normalization constant such that $\sum_j w(x_{ij}, x_{kl}) = 1$, $S(N_{ij})$ and $S(N_{kl})$ are vectors containing the spectra of neighborhoods N_{ij} and N_{kl} of voxels x_{ij} and x_{kl} respectively and h serves as a smoothing parameter [5].

To increase the robustness of our method for MRSI data, in the next section we propose a dephasing approach tailored for use in the frequency-phase NLM.

2.2 Spectral Dephasing

For the acquired data I , as the spectral phase $\Phi(I)$ is unknown, the probability of finding a similar neighborhood spectra are very low. To counter this effect, a dephasing step is performed to consider a wide range of possible phase variations in the pattern analysis. For each voxel x_{ij} , the complex time-domain signal $S_{ij}(t)$ is shifted by a set of phase angles Θ . This is given by $S_{ij}^\Theta(t) = S_{ij}(t) \cdot e^{-i\Theta}$, where $\Theta \in [-n_1\pi, (n_2 + 2)\pi]$, $\{n_1, n_2 \in \mathbb{R} \mid n_1, n_2 \geq 0\}$. Θ here is defined to be the range of angles through which the spectrum can be shifted. The dephased signal is transformed into the frequency-domain, $S_{ij}^\Theta(t) \xrightarrow{\mathcal{F}} S_{ij}^\Theta(\omega)$, following which its real component, $\mathbb{R}(S_{ij}^\Theta(\omega))$, is taken to generate a 2D spectral-phase matrix. Note that in this 2D matrix generated, for each voxel x_{ij} , the imaginary part at a given Θ is $\mathbb{I}(\Theta) = \mathbb{R}(\Theta + \pi/2)$. This approach is illustrated in Fig. 1.

Repeating this step for all MN voxels gives us a 3-D dataset on which the NLM is implemented to give the denoised spectra $\hat{S}_{ij}^\Theta(\omega) \in \mathbb{R}$. Our approach has 2 key innovations: the denoising method is (i) robust to phase shifts as the range of angles considered varies from 0 to 2π periodically for all spectral signals, and is (ii) adaptive to the imaging sequence as the spectrum is denoised by relying on similar signals in the given data and not on predefined prior assumptions.

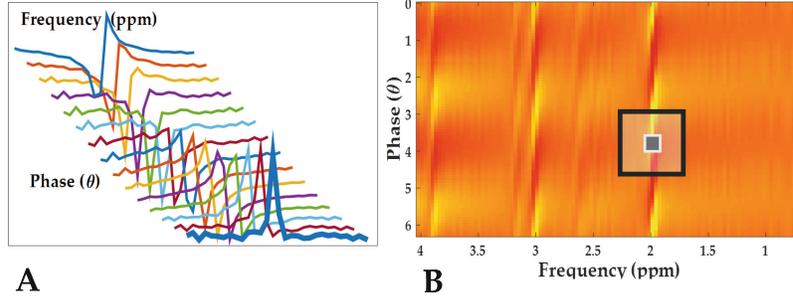


Fig. 1. MRSI data dephasing shown here for a sample voxel: (A) Changes in spectral pattern as it is shifted by different phase angles. (B) Corresponding 2D frequency-phase image space generated for the voxel. A sample patch (black box) is selected and then denoised by the NLM-based matching in the frequency-phase space.

Algorithm 1. Frequency-Phase NLM denoising for MRSI

MRSI Input $I : \Omega^3 = (S(\omega) \times M \times N) \mapsto \mathbb{C}$

Define Phase Angle range $\Theta : [-n_1\pi : (n_2 + 1)\pi] \{n_1, n_2 \in \mathbb{R} \mid n_1, n_2 \geq 0\}$

Spectral Dephasing:

$S_{ij}(t) = \mathcal{F}^{-1}(S_{ij}(\omega))$: for each voxel $ij \in (M \times N)$

$S_{ij}^\Theta(\omega) = \mathcal{F}(S_{ij}(t) \cdot e^{-i\Theta})$: Phase shift by Θ

NLM: $\hat{S}_{ij}^\Theta(\omega) = NLM[\mathbb{R}(S_{ij}^\Theta(\omega)) \times MN] : \forall$ voxels

Spectral Rephasing:

$\mathbb{C}(\hat{S}_{ij}^\Theta(t)) = \mathcal{F}^{-1}[\mathbb{R}(\hat{S}_{ij}^\Theta(\omega)) + \mathbb{I}(\hat{S}_{ij}^\Theta(\omega))]$: Re-generate complex data for all Θ

$\mathbb{C}(\hat{S}_{ij}^{-\Theta}(\omega)) = \mathcal{F}(\mathbb{C}(\hat{S}_{ij}^\Theta(t)) \cdot e^{i\Theta})$: Rephase by $-\Theta$

Frequency-Phase NLM Output: $\mathbb{C}(\hat{S}_{ij}(\omega)) = \text{mean}(\mathbb{C}(\hat{S}_{ij}^{-\Theta}(\omega))) \forall \Theta$

2.3 Spectral Rephasing and Recombination

Post-NLM, $\hat{S}_{ij}^\Theta(\omega)$ is rephased in order to generate the denoised complex signal $\mathbb{C}(\hat{S}_{ij}(\omega))$. The complex spectral signal $\mathbb{C}(\hat{S}_{ij}^\Theta(\omega))$ is re-generated $\forall \Theta$ by combining $\mathbb{R}(\hat{S}_{ij}^\Theta(\omega))$ and $\mathbb{I}(\hat{S}_{ij}^\Theta(\omega)) (= \mathbb{R}(\hat{S}_{ij}^{\Theta+\pi/2}(\omega)))$. The equivalent time signal is obtained by $\mathbb{C}(\hat{S}_{ij}^\Theta(\omega)) \xrightarrow{\mathcal{F}^{-1}} \mathbb{C}(\hat{S}_{ij}^\Theta(t))$. After this, $\mathbb{C}(\hat{S}_{ij}^\Theta(t))$ undergoes an inverse phase shift by $-\Theta$ to remove the dephasing effect as given by $\mathbb{C}(\hat{S}_{ij}^{-\Theta}(t)) = \mathbb{C}(\hat{S}_{ij}^\Theta(t)) \cdot e^{i\Theta}$. This re-phased signal is transformed back to the spectral domain to obtain $\mathbb{C}(\hat{S}_{ij}^{-\Theta}(\omega))$. Thereafter, the $\mathbb{C}(\hat{S}_{ij}^{-\Theta}(\omega))$ are averaged over all Θ to generate a single complex spectra $\mathbb{C}(\hat{S}_{ij}(\omega))$. The entire pipeline for dephasing and rephasing the spectra is shown in Algorithm 1.

3 Experiments and Results

We performed two different experiments to test the improvement in SNR and metabolite quantification using our proposed denoising method. In the first experiment, we evaluate our method on the publicly available BrainWeb database [4],

while in the second experiment we use human in-vivo MRSI data. The SNR of a metabolite was calculated by dividing the maximum value of the metabolite peak by the standard deviation of the spectral region having pure noise. For both experiments, we tested with different noise levels against a ground-truth data to assess the improvement in SNR and spatial-spectral resolution.

3.1 Data Acquisition

We used BrainWeb to simulate a brain MRSI image (size: 64×64 voxels, slice thickness = 1 mm, noise level = 3%) with segmented tissue types, namely White Matter (WM), Grey Matter (GM) and Cerebro Spinal Fluid (CSF) as shown in Fig. 2. In order to have a comparable spectrum with the in-vivo data, water was added to the signal and the main metabolites- NAA, Cho and Cr- were simulated using Priorset (Vespa) [1]. Metabolite concentrations for WM, GM and CSF were based on commonly reported literature values [10,14]. Next, Gaussian noise of levels 2 and 3 times the standard deviation, σ , of the original image data were added to the ground-truth signal.

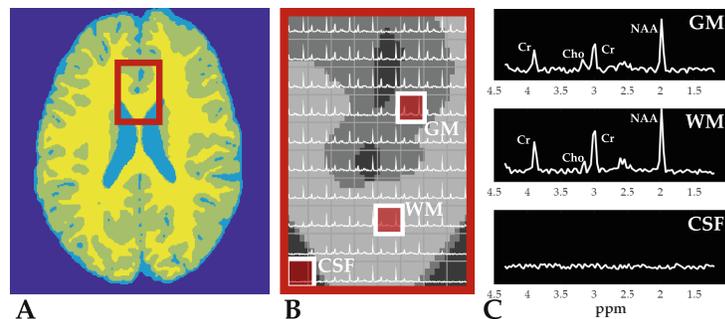


Fig. 2. Simulated brain MRSI dataset. (A) The simulated brain with the region of interest (red box). (B) Highlighted regions corresponding to GM, WM and CSF (c) Corresponding spectrum of GM, WM and CSF. Note that CSF has only water. (Color figure online)

In the case of in-vivo data, we acquired a 2D-MRSI data of the brain of a healthy human volunteer using a 3T-HDxt system (GE-Healthcare). PRESS localization [2], CHES water suppression [6] and EPSI readout [9] were used as part of the sequence. The acquisition parameters were: Field of view (FOV) = $160 \times 160 \times 10 \text{ mm}^3$, voxel size = $10 \times 10 \times 10 \text{ mm}^3$, TE/TR=35/2000 ms and spectral bandwidth = 1 kHz. The dataset was zero-filled and reconstructed to generate a grid of 32×32 voxels and 256 spectral points. 6 (**ground truth**), 3 and 1 averages were acquired with a total scan duration of 33 min (5.5 min per average). Figure 4(A) shows the in-vivo data acquired along with the entire field-of-view (white grid) and the corresponding spectra of a voxel (red box).

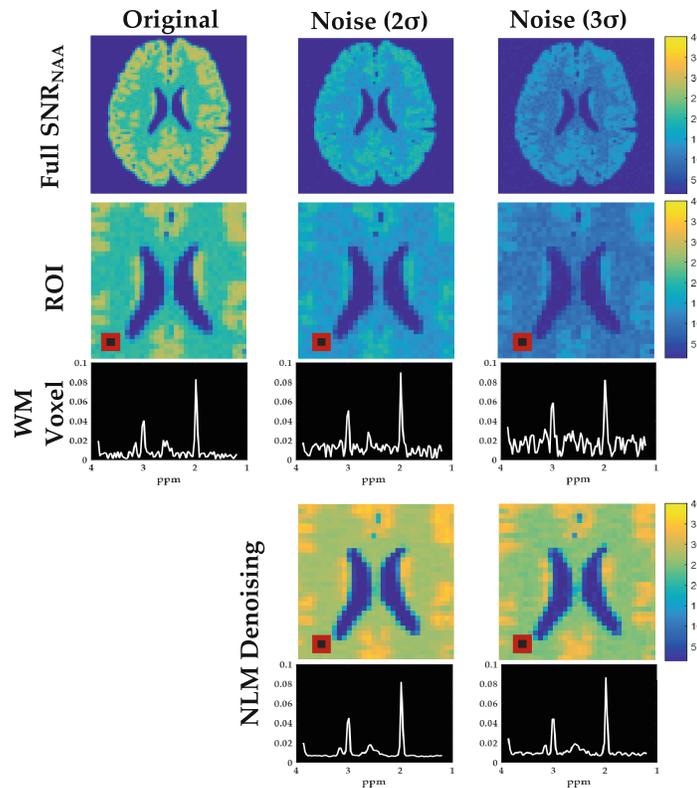


Fig. 3. NLM Denoising results in the simulated data. (From Top) Row 1 (L-R): Full SNR of NAA in – original data, with additive noise 2σ , and noise 3σ (σ is the standard deviation of the original data). Row 2 & 3: 32×32 Region of Interest (ROI) for applying the frequency-phase NLM: SNR of NAA in the original data, noise level 2σ , 3σ and the corresponding spectra of reference WM voxel (red box). Row 4 & 5 (L-R): Denoised SNR for noise level 2σ (SNR improvement = 2.9), for noise level 3σ (SNR improvement = 2.2), and the corresponding denoised spectrum. The SNR improves significantly while retaining the spatial-spectral resolution (seen by no voxel bleeding in the CSF). (Color figure online)

3.2 Results

Simulated data: In Fig. 3, we show the SNR improvement for NAA for data with noise levels 2σ and 3σ in a 32×32 region of interest. It is evident that while the spectral SNR improves significantly, the spatial resolution is preserved as the lower concentration metabolite peaks have only a small amount of smoothing and there is no voxel bleeding in the CSF (containing only water).

In-vivo data: Figure 4 reports the SNR improvement in NAA for the 3-averages and the 1-average data as compared to the ground-truth 6-averages data. The figure also presents the results from the LCMoel [11] which is the gold standard quantitation tool in MRS analysis. LCMoel fits the spectral signal $S(\omega)$ using a basis set of spectra of metabolites acquired under identical acquisition conditions as the in-vivo data. As explained earlier for noisy data, the non-linear fitting leads to poor spatial resolution. Therefore, the LCMoel can be used to

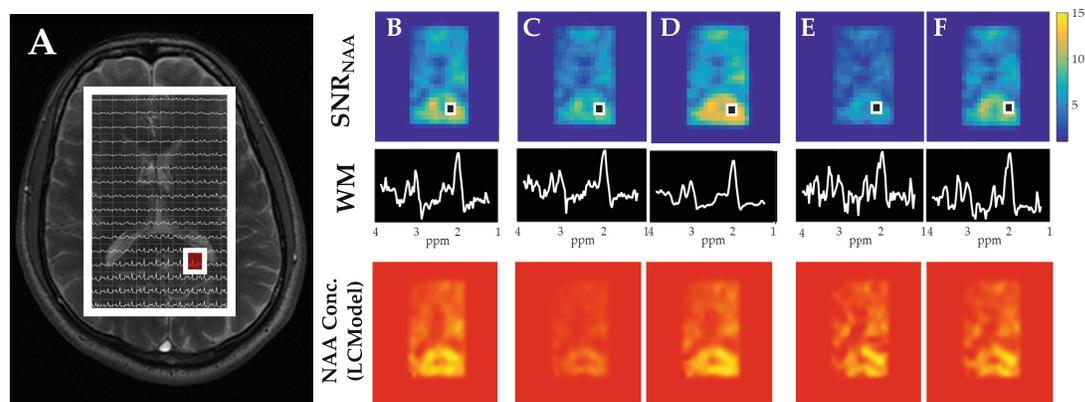


Fig. 4. Denoising results for in-vivo data. (A) Original human in-vivo brain MRSI data with the excitation region shown (white grid). (From Top) Row 1 & 2: SNR of NAA and the corresponding WM voxel spectra in: (B) 6-averages data (**ground-truth**), (C) 3-averages data (original) and (D) denoised, (E) single-scan data (original) and (F) denoised with corresponding spectra. Row 3: LCModel based absolute concentration estimate of NAA in: 6-averages data, 3-averages data (original) and denoised, 1-average data (original) and denoised. The NAA concentration estimate and spectral SNR improve considerably as seen in columns D (SNR = 23.29) and F (SNR = 11.38) against the ground-truth (SNR = 11.44).

Table 1. SNR and LCM Quantification results for NAA, Cho and Cr before and after using frequency-phase NLM on in-vivo MRSI data. Mean SNR for the denoised data is comparable or better than the ground-truth data while the FWHM of the water peak is lower than the ground-truth data thereby preserving spatial-spectral resolution.

Data	Mean SNR (NAA)	SNR improvement (NAA)	Mean SNR (Cho)	SNR (Cho)	Mean SNR (Cr)	SNR improvement (Cr)	FWHM
1-average	5.70	–	3.74	–	4.04	–	0.125
3-averages	8.82	–	5.59	–	5.97	–	0.135
6-averages (ground-truth)	11.44	–	6.96	–	7.61	–	0.135
Single scan (NLM)	11.38	1.98	6.89	1.82	7.78	1.90	0.130
3-averages (NLM)	23.29	2.63	14.08	2.48	15.59	2.57	0.134

assess the improvement in spatial resolution through a better fit. Due to space constraints, we present the results for NAA only and mention the SNR values for Cho and Cr. LCModel quantification (Fig. 4) shows that the absolute concentration estimation of NAA in the denoised data improves significantly. The Full-Width Half Maximum (FWHM) shows information about the water peak – a narrow peak gives a better spatial resolution. As shown in Table 1, the FWHM of the denoised 1- and 3-averages data is lower than the 6-averages data while the corresponding mean SNR improves considerably. Therefore, we observe here that our method can accelerate MRSI data acquisition by almost 2 times by reducing the number of scans acquired.

4 Conclusion

In this work, we proposed a novel frequency-phase NLM-based denoising method for MRS Imaging to improve the SNR and spatial resolution of the metabolites. A spectral dephasing approach is promoted to compensate for the unknown phase information of the acquired data. To the best of our knowledge, this is a novel application of the concept of NLM and has been validated on both simulated and in-vivo MRSI data.

In particular, we assessed the effect of our method on metabolites such as NAA, Cho and Cr and obtained a visible improvement in SNR while the spatial resolution was preserved which, subsequently, led to a better estimation of the absolute concentration distribution of NAA. This has direct benefits as it would accelerate data acquisition by taking fewer scan averages. Future work would involve using a more robust metabolite-specific search in the given dataset with less smoothing. This can be coupled with optimal computational efficiency and better estimation of the in-vivo metabolites.

Acknowledgments. The research leading to these results has received funding from the European Union's H2020 Framework Programme (H2020-MSCA-ITN-2014) under grant agreement no 642685 MacSeNet.

References

1. Vespa project (Versatile simulation, pulses and analysis). <https://scion.duhs.duke.edu/vespa/project>
2. Bottomley, P.A.: Spatial localization in NMR spectroscopy in vivo. *Ann. N. Y. Acad. Sci.* **508**, 333–348 (1987). doi:[10.1111/j.1749-6632.1987.tb32915.x](https://doi.org/10.1111/j.1749-6632.1987.tb32915.x)
3. Buades, A., Coll, B.: A non-local algorithm for image denoising. *Comput. Vis. Pattern* **2**(0), 60–65 (2005)
4. Collins, D.L., Zijdenbos, P., Kollokian, V., Sled, J.G., Kabani, N.J., Holmes, C.J., Evans, C.: Design and construction of a realistic digital brain phantom. *IEEE Trans. Med. Imaging* **17**(3), 463–468 (1998)
5. Coupé, P., Yger, P., Prima, S., Hellier, P., Kervrann, C.: An optimized blockwise non local means denoising filter for 3D magnetic resonance images. *IEEE Trans. Med. Imaging* **27**(4), 425–441 (2008)
6. Haase, A., Frahm, J., Hänicke, W., Matthaei, D.: 1H NMR chemical shift selective (CHESS) imaging. *Phys. Med. Biol.* **30**(4), 341–344 (1985). <http://stacks.iop.org/0031-9155/30/i=4/a=008>
7. Kelm, B.M., Kaster, F.O., Henning, A., Weber, M.A., Bachert, P., Boesiger, P., Hamprecht, F.A., Menze, B.H.: Using spatial prior knowledge in the spectral fitting of MRS images. *NMR Biomed.* **25**(1), 1–13 (2012)
8. Nguyen, H.M., Peng, X., Do, M.N., Liang, Z.: Spatiotemporal denoising of MR spectroscopic imaging data by low-rank approximations. *IEEE ISBI: From Nano to Macro* **0**(3), 857–860 (2011)
9. Posse, S., DeCarli, C., Le Bihan, D.: Three-dimensional echo-planar MR spectroscopic imaging at short echo times in the human brain. *Radiology* **192**(3), 733–738 (1994). <http://pubs.rsna.org/doi/abs/10.1148/radiology.192.3.8058941>

10. Pouwels, P.J.W., Frahm, T.: Regional metabolite concentrations in human brain as determined by quantitative localized proton MRS. *Magn. Reson. Med.* **39**(1), 53–60 (1998)
11. Provencher, S.W.: Estimation of metabolite concentrations from localized in vivo proton NMR spectra. *Magn. Reson. Med.* **30**(6), 672–679 (1993). <http://www.ncbi.nlm.nih.gov/pubmed/8139448>
12. de Graaf, R.A.: *In Vivo NMR Spectroscopy: Principles and Techniques*, 2nd edn. Wiley, Hoboken (2013)
13. Schulte, R.F., Lange, T., Beck, J., Meier, D., Boesiger, P.: Improved two-dimensional J-resolved spectroscopy. *NMR Biomed.* **19**(2), 264–270 (2006)
14. Wang, Y., Li, S.: Differentiation of metabolic concentrations between gray matter and white matter of human brain by in vivo ¹H magnetic resonance spectroscopy. *Magn. Reson. Med.* **39**(1), 28–33 (2005)

Quantification of Metabolites in Magnetic Resonance Spectroscopic Imaging Using Machine Learning

Dhritiman Das^{1,3} (✉), Eduardo Coello^{2,3}, Rolf F. Schulte³,
and Bjoern H. Menze¹

¹ Department of Computer Science,
Technical University of Munich, Munich, Germany
`dhritiman.das@tum.de`

² Department of Physics, Technical University of Munich, Munich, Germany

³ GE Global Research Europe, Munich, Germany

Abstract. Magnetic Resonance Spectroscopic Imaging (MRSI) is a clinical imaging modality for measuring tissue metabolite levels in-vivo. An accurate estimation of spectral parameters allows for better assessment of spectral quality and metabolite concentration levels. The current gold standard quantification method is the LCModel - a commercial fitting tool. However, this fails for spectra having poor signal-to-noise ratio (SNR) or a large number of artifacts. This paper introduces a framework based on random forest regression for accurate estimation of the output parameters of a model based analysis of MR spectroscopy data. The goal of our proposed framework is to learn the spectral features from a training set comprising of different variations of both simulated and in-vivo brain spectra and then use this learning for the subsequent metabolite quantification. Experiments involve training and testing on simulated and in-vivo human brain spectra. We estimate parameters such as concentration of metabolites and compare our results with that from the LCModel.

1 Introduction

Magnetic resonance spectroscopic imaging (MRSI) is an in-vivo clinical imaging modality which detects nuclear magnetic resonance signals produced by nuclei in living tissues. Quantification of this signal amplitude generates metabolic maps which show the concentration of metabolites in the sample being investigated. Accurate quantification of these metabolites is important for diagnosis of brain tumor and other in-vivo diseases. For this purpose, a common practice in the MRS community has been to use non-linear spectral fitting tools such as the LCModel [5], TARQUIN [9], AMARES [8] and ProFit [7] amongst which the LCModel is regarded as the gold standard fitting tool. In this study, we present an alternative to the non-linear model fitting using a machine learning approach.

Non-linear Model Fitting. The LCModel software uses a linear combination of metabolite basis spectra set to model the spectral measurement in the fre-

quency domain. It also uses smoothing splines to model the baseline signals and subsequently fits the parameters of the basis set using a non-linear optimisation. LCModel incorporates the prior knowledge of the data while modeling the fit and this ensures robustness in the model leading to estimation of the spectral parameters such as concentration of metabolites. Some of the drawbacks of this non-linear fitting model are: (1) Metabolite quantification can be time-consuming depending on the dataset size and requires a lot of manual parameter tuning. (2) The error in estimating parameters is lower if high SNR spectra are used since the non-linear voxel-wise fitting to noisy data leads to a high amount of local minima and subsequent inaccuracy in quantification [3,4].

Machine Learning. Machine learning methods such as decision forests, random forests [2] are being extensively used in the medical imaging community for tasks such as parameter estimation, diseases diagnosis, segmentation, etc. In MRSI, machine learning tools have been used only for specific tasks such as classification of spectra [4] and assessment of spectral quality [1]. This opens up the possibility of using the recent advances in machine learning to predict MRSI data parameters while addressing the drawbacks of conventional fitting tools such as long computation time and poor performance for data with artifacts.

Our Contribution. In this work, we propose a simple yet effective method using random forest regression for multi-parameter estimation in MR Spectroscopic Imaging. We generate over 1 million simulated spectra training-set having concentration magnitudes, linewidth effects, baseline and lipid artifacts. We also use spectral data from 287 human subjects to create a physical training model to be used in the regression framework (Sect. 3.1). In the following we present our method adapting random forest regression to MRSI (Sect. 2) followed by experiments in the aforementioned dataset. Our proposed method is then validated quantitatively and qualitatively using: (1) synthetic brain spectra, (2) human in-vivo single voxel spectra having the same image acquisition protocol as the physical training model and (3) independently acquired human in-vivo 2D MRS Images to perform a blind test on the physical and synthetic models. We present the results (Subsect. 3.2) of our experiments followed by a summary and discussion (Sect. 4) on the future work in this domain. This is the first application- to the best of our knowledge- of machine learning for determining MRS parameters which were otherwise determined using basis fitting tools.

2 Methods

MR Spectroscopy. Magnetic resonance spectroscopy, based on the concept of nuclear magnetic resonance (NMR), exploits the resonance frequency of a molecule, to obtain information about the concentration of a particular metabolite [6]. The time-domain complex signal of a nuclei is given by:

$$S(t) = \int p(\omega) \exp(-i\Phi) \exp(-t/T_2^*) dw. \quad (1)$$

The frequency-domain signal is given by $S(\omega)$, T_2^* is the magnetization decay in the transverse plane due to magnetic field inhomogeneity and $p(\omega)$ comprises of Lorentzian absorption and dispersion line-shapes function having the spectroscopic information about the sample. Φ represents the phase, $(\omega t + \omega_0)$, of the acquired signal where ωt is the time-varying phase change and ω_0 is the initial phase. Non-linear fitting tools facilitate the generation of metabolic maps to estimate concentration of metabolites such as N-acetyl-aspartate (NAA), Creatine (Cr) and Choline (Cho). An example of the spectra present in the brain has been shown in Fig. 1.

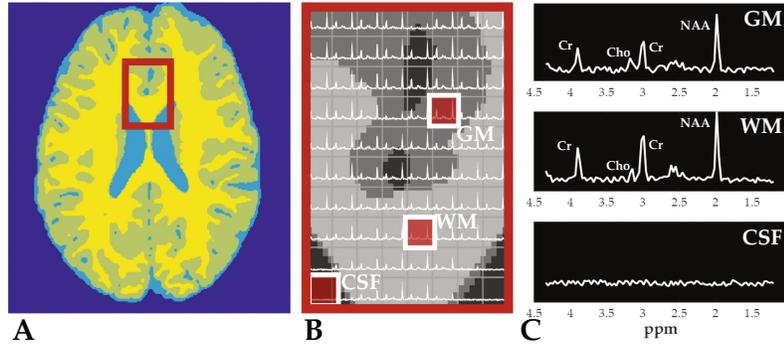


Fig. 1. Example brain 2D MRSI dataset. (A) The simulated brain with the region of interest (red box). (B) Highlighted regions corresponding to GM, WM and CSF (c) Corresponding spectrum of GM, WM and CSF

Random Forest Regression. Random Forests [2] have been shown to be effective in a wide range of classification and regression problems. These comprise of a set of binary trees wherein splits are created in each tree based on a random subsets of the feature variables on which the forests are subsequently trained. Piecewise linear regression is implemented by each tree over the input data and, after seeking for the best prediction at every node, data points are sent to the left or right branches based on feature selection by thresholding. This process continues till it reaches the end of the tree and subsequently the weighted average of the prediction from each tree is taken to give a single output estimate. The randomness in the training process encourages the trees to give independent estimates which can be combined to achieve an accurate and robust result.

For MRSI, we adapt the random forest approach to have a training dataset $D = (S_i(\omega), Y_i)$, $i \in [1, N]$, where N is the total number of training spectra. $S_i(\omega)$ represents the training spectral data while Y_i represents the corresponding multi-parameter training labels. For our model, we consider the concentrations of NAA, Cho and Cr for simulated data, while for the real data we additionally consider Myo-Inositol (mI) and Glutamate+Glutamine (Glx). Therefore, for a given spectra $S_i(\omega)$, $Y_i = [NAA_i, Cho_i, Cr_i, mI_i, Glx_i]$.

Running the random forest regression on this produces a training model which can then be used to obtain parameter estimates \hat{Y}_j of test spectra $S_j(\omega)$ having test labels Y_j , $j \in [1, M]$ where M is the total number of test spectra.

Error Calculation. For our experiments, given the estimate \hat{Y}_j and the testing label Y_j , the estimate error for the parameter Y_j can be calculated as,

$$\hat{E}_j = \|\hat{Y}_j - Y_j\| / \|Y_j\| \quad (2)$$

This method helps us to assess the change in parameter estimate over the testing/ground-truth values.

3 Experiments and Results

3.1 Data

We perform 4 sets of experiments to assess our proposed method: (1) training and testing on simulated spectra (**Synthetic - Synthetic (Spectra)**), (2) training and testing on human in-vivo spectral data from different subjects but having the same acquisition protocol (**Real (Spectra) - Real (Spectra)**), (3) training and testing on human in-vivo spectral data from different subjects with different acquisition protocol (**Real (Spectra) - Real (MRS Images)**) and (4) using the simulated spectra model to test on MRS images (**Synthetic (Spectra) - Real (MRS Images)**).

Synthetic (Spectra). A metabolite basis set was generated by using the data provided by the ISMRM MRS Fitting Challenge 2016. These were then used to simulate over 1 million spectra. In order to ensure that the simulated spectra was as close as possible to human in-vivo spectra, we incorporate the following features: variations in NAA, Cho, Cr concentrations, macro-molecular baseline, lipids, t2 values (for changes in linewidth) and signal-to-noise ratio (SNR) to account for changes in spectral quality. As a preliminary case study, we only simulate the major metabolites (NAA, Cho and Cr) as these are easily detected by the LCModel and would, therefore, help us to evaluate the outcome of our approach and allow a suitable comparison with the LCModel. A set of over 10,000 independent test spectra were also simulated with varying combinations of the aforementioned features. For both the training and testing sets, we used the basis-set metabolite concentration values as our ground-truth.

Real (Spectra). To evaluate our method on in-vivo data, we utilize LCModel-fitted single-voxel spectroscopy (SVS) data from 287 independent human subjects. The data was obtained using the same standardized imaging protocol with the following acquisition parameters: TE/TR = 35/2000 ms, spectral width = 2500 Hz, number of points = 1024. We implement a K-fold cross-validation with 10 folds along with the random-forest regression to generate different training and testing sets having spectra from 259 and 28 subjects respectively. The metabolites assessed were: NAA, Cho, mI and Glx.

Real (MRS Images). To further assess our approach, we acquire a standard phase-encoded 2D brain MRSI data of a healthy human volunteer on a 3T scanner using a point-resolved spin-echo localization sequence (PRESS) with

voxel size = $10 \times 10 \times 15$ mm³, TE/TR = 35/1000 ms, spectral width = 2000 Hz, number of points = 400. For testing purposes, we use 96 spectra from the inner-region of the brain which serves as the region of interest.

Due to the differences in acquisition parameters of the training and testing set, both the resulting spectra vary in amplitude and metabolite peak alignment. We perform a pre-processing spectral alignment step where all the test spectra are cropped from 4.3 to 0.2 ppm and interpolated to the same number of points as the training spectra to compensate for differences in acquisition bandwidth. This is followed by normalizing the amplitude of the test spectra using one of the training spectra as reference.

3.2 Results

Synthetic - Synthetic (Spectra). We perform an initial experiment to determine the out-of-bag (OOB) error using different number of trees and features on a set of 20,000 simulated train and test spectra. Based on the results shown in Fig. 2, we proceed with the parameter estimation experiment by identifying the appropriate number of trees and features required to achieve convergence of the OOB error. For the regression error estimates, we use metabolite concentration ratios with respect to Cr (used as a standard assessment method in MRS as a means for calibration). We obtain R scores of 0.968 and 0.962 for NAA/Cr and Cho/Cr values respectively. The corresponding figures representing the linear regression are shown in Fig. 3 and the error plots in comparison with the LCMoel are shown in Fig. 4.

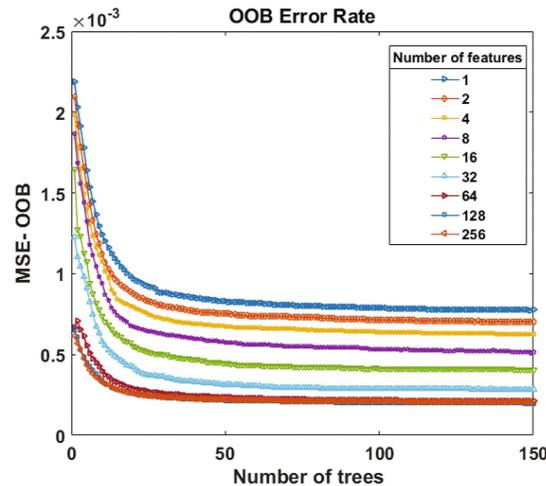


Fig. 2. Out-Of-Bag (OOB) Error for Simulated Spectra. The experiment is performed for a varying number of features (from 1 to 256 as shown in the legend) and each iteration is assessed for a varying number of trees (as shown in the X-axis). The Y-axis represents the OOB Error rate. The error rate is minimal for more than 64 features and also converges when the number of trees is close to 100.

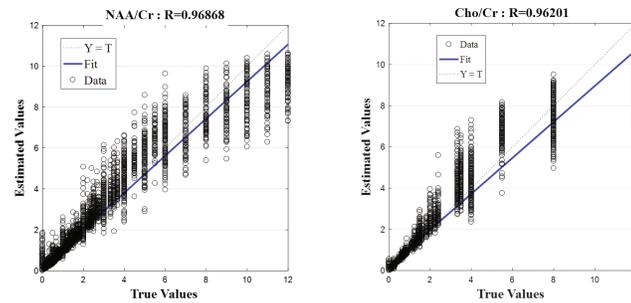


Fig. 3. Regression Scores for the following parameters (from left to right): NAA/Cr concentration estimate and Cho/Cr concentration estimate. The X-axis represents the true values of the parameter while the y-axis represents the estimated values. Both sets of values are plotted using linear regression.

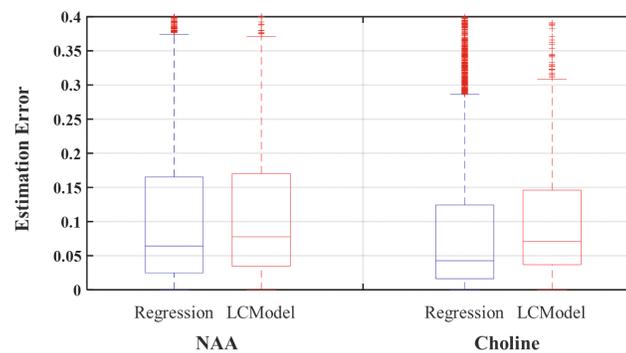


Fig. 4. Synthetic-Synthetic (Spectra): Estimation error for different metabolite concentration ratios in a given test-set. Whiskers span the [min max] values. Median error values are represented by the red line and are as follows: **NAA/Cr** Regression = 0.064, LCModel = 0.077, **Cho/Cr** Regression = 0.043, LCModel = 0.070.

Real (Spectra) - Real (Spectra). For the SVS dataset, we use the LCModel concentration ratio estimates as the ground-truth. Table 1 indicates the mean metabolite concentration estimate error across the 10-folds of the cross-validation process using the random forest regression method. Median error for the NAA/Cr estimate is 0.068, 0.072 for the Cho/Cr estimate, 0.093 for the mI/Cr estimate and 0.070 for the Glx/Cr estimate compared to the corresponding LCModel estimates. The difference in error estimates is small and shows a similarity in assessment between our proposed method and the LCModel. Moreover, the low-concentration metabolites such as mI and Glx usually display a fitting error with the LCModel and the estimation error for these metabolite ratio concentrations is lower indicating that our model works well for these metabolites as well.

Synthetic (Spectra) - Real (Images). We test our synthetic spectra training model on the 2D MRSI data and the results are shown in the boxplot in Fig. 5 along with the resulting concentration distribution from both the regression approach and the non-linear model fit. As our synthetic model is trained for only

Table 1. Concentration-ratio estimate errors using random forest regression. Results are for the experiments **Real(spectra)-Real(spectra)** and **Real(spectra)-Real(Images)**. The errors are calculated over the respective LCMoel estimates as per the formula given in Eq. 2. The major metabolites (**NAA** and **Cr**) show a low error while the smaller concentration metabolites (**mI** and **Glx**) show a slightly higher error.

	Naa/Cr	Cho/Cr	mI/Cr	Glx/Cr
Real-Real (Spectra)	0.068	0.072	0.093	0.070
Real-Real (Images)	0.1	0.18	0.217	0.13

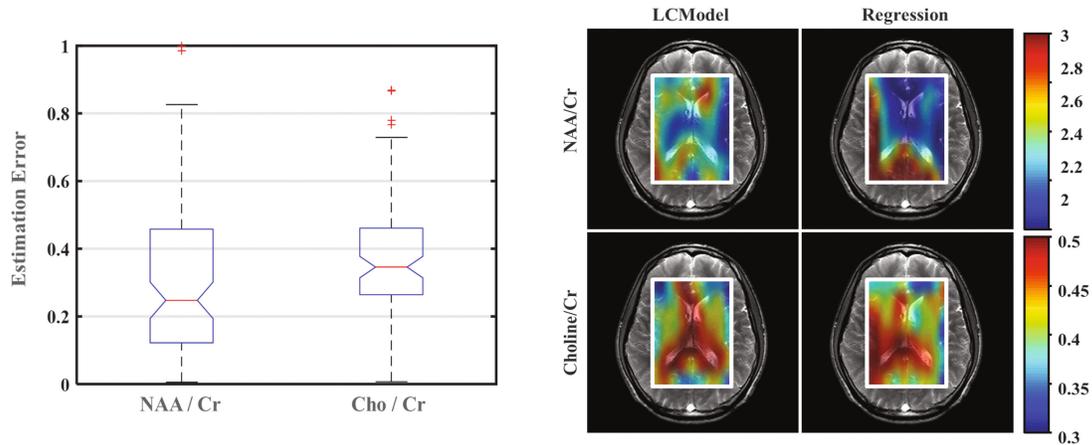


Fig. 5. **Left:** Synthetic (Spectra)-Real (MRS Images): Estimation error for different metabolite concentration ratios for the same test dataset. Whiskers span the [min max] values. Median error values are represented by the red line and are as follows: NAA/Cr = 0.024, Cho/Cr = 0.034. **Right:** NAA/Cr and Cho/Cr concentration distribution estimates from random forest regression and non-linear model fit.

NAA and Cho ratios, we show the errors for these two only. Median estimate error for NAA/Cr is 0.24 using regression. For Cho/Cr, the estimation error is 0.34. The corresponding concentration values estimated from the LCMoel serves as our ground-truth.

Real (Spectra) - Real (Images). We perform a blind test with 96 2D MRSI spectra against the training model generated using the 287 SVS spectra and the results are shown in Table. 1. Median estimate error for NAA/Cr is 0.1, for Cho/Cr is 0.18, for mI/Cr is 0.217 and for Glx/Cr is 0.13. Although we expect the errors to be higher in the blind test due to difference in the acquisition protocols of the training and testing dataset, the errors appear to be within a reasonable window. As expected, the estimated errors are highest for mI/Cr while Glx/Cr surprisingly has a lower error than Cho/Cr.

The Real Spectra training model provides a marginally better metabolite concentration estimate than the Synthetic spectra model. We attribute this to the presence of arbitrary scanning effects and artifacts in the real spectra model

as compared to the synthetic model. For future experiments, this provides the scope for learning on a large synthetic spectral data-set with similar additional arbitrary effects to have a robust classifier for real data (especially in the cases where annotating training data is expensive).

4 Conclusion

Machine learning techniques such as Random Forest-based regression provide a new and faster way of metabolite quantification. Our synthetic training model accounts for spectral features such as macro-molecular baseline, lipids, linewidth and SNR variations in combination with different metabolite concentrations. Additional features such as frequency and/or phase-shift effects along with B0 inhomogeneity could be incorporated in the model to improve robustness. For the human in-vivo data, we use training spectra from different subjects and the random-forest regression provides a low amount of estimation error over the LCModel fit even in the presence of arbitrary scanning effects. Training times for the simulated spectra can be considerable (around 5–6 h) given that we generate over 1 million spectra while it is only a few minutes for the in-vivo spectra. On the other hand, testing and concentration estimation happens in only a few seconds and is considerably faster than the non-linear model fitting. The machine learning approach may be used directly, or indirectly by initializing LCModel fits thereby improving their results in the presence of noise and speeding up convergence. They can also be combined with global decisions about spectral quality predicting whether a spectrum can or cannot be interpreted by the physics model because of the presence of artifacts.

Future work would involve using a more robust approach such as deep-learning based methods to improve the accuracy of parameter estimation. Once a framework has been established, further work can be done on having disease-based training models for parameter estimation to predict disease progression and the corresponding metabolite maps.

References

1. Pedrosa de Barros, N., McKinley, R., Wiest, R., Slotboom, J.: Improving labeling efficiency in automatic quality control of MRSI data. *Magnetic Resonance in Medicine* (2017). <http://dx.doi.org/10.1002/mrm.26618>
2. Breiman, L.: Random forests. *Mach. Learn.* **45**(1), 5–32 (2001). <http://dx.doi.org/10.1023/A:1010933404324>
3. Kelm, B.M., Kaster, F.O., Henning, A., Weber, M.A., Bachert, P., Boesiger, P., Hamprecht, F.A., Menze, B.H.: Using spatial prior knowledge in the spectral fitting of MRS images. *NMR Biomed.* **25**(1), 1–13 (2012)
4. Menze, B.H., Kelm, B.M., Weber, M.A., Bachert, P., Hamprecht, F.A.: Mimicking the human expert: pattern recognition for an automated assessment of data quality in MR spectroscopic images. *Magn. Reson. Med.* **59**(6), 1457–1466 (2008). <http://dx.doi.org/10.1002/mrm.21519>

5. Provencher, S.W.: Estimation of metabolite concentrations from localized in vivo proton NMR spectra. *Magn. Reson. Med.* **6**, 672–679 (1993)
6. de Graaf, R.A.: *In Vivo NMR Spectroscopy: Principles and Techniques*, 2nd edn. Wiley, Chichester (2013)
7. Schulte, R.F., Lange, T., Beck, J., Meier, D., Boesiger, P.: Improved two-dimensional J-resolved spectroscopy. *NMR Biomed.* **19**(2), 264–270 (2006)
8. Vanhamme, L., van den Boogaart, A., Huffel, S.V.: Improved method for accurate and efficient quantification of MRS data with use of prior knowledge. *J. Magn. Reson.* **129**(1), 35–43 (1997). <http://www.sciencedirect.com/science/article/pii/S1090780797912441>
9. Wilson, M., Reynolds, G., Kauppinen, R.A., Arvanitis, T.N., Peet, A.C.: A constrained least squares approach to the automated quantitation of in vivo 1h magnetic resonance spectroscopy data. *Magn. Reson. Med.* **65**(1), 1–12 (2011). <https://dx.doi.org/10.1002/mrm.22579>

Direct Estimation of Model Parameters in MR Spectroscopic Imaging using Deep Neural Networks

Dhritiman Das^{1,4}, Eduardo Coello^{2,4}, Anjany Sekuboyina^{1,3}, Rolf F. Schulte⁴,
Bjoern H. Menze¹

¹Department of Computer Science, Technical University of Munich, Germany

²Department of Physics, Technical University of Munich, Germany

³Klinikum rechts der Isar, Munich, Germany ⁴GE Healthcare, Munich, Germany

Abstract. We introduce a deep-learning based framework based on a multilayer perceptron for estimation of the output parameters of a model-based analysis of MR spectroscopy data. Our proposed framework: (1) learns the spectral features from a training set comprising of different variations of synthetic spectra; (2) uses this learning and performs non-linear regression for the subsequent metabolite quantification. Experiments involve training and testing on simulated and in-vivo human brain spectra. We estimate parameters such as metabolite-concentration ratios and compare our results with that from the LCModel.

1 Introduction

Quantification of MR Spectroscopy (MRS) signals generates metabolic maps which show the concentration of metabolites in the sample being investigated. Accurate quantification of these metabolites is important for diagnosis of brain tumor and other in-vivo diseases. For this purpose, non-linear model-fitting tools are widely used (such as the LCModel [4], TARQUIN [8], AMARES [6] and ProFit [5]). The LCModel is widely regarded as the gold-standard fitting tool. However, some of its drawbacks include: (1) prior knowledge-tuning and long fitting times, and (2) high estimation error for noisy data. Prior work has also focused on using machine-learning for metabolite-quantification [2]. In this study, we present an alternative to the non-linear model fitting using a deep-learning approach.

2 Methods

A multilayer perceptron(MLP)[7] is a fully-connected, feedforward deep-neural network comprising of three or more layers of non-linearly activated nodes. The nodes in each layer are connected to the next layer with certain weights and a supervised learning technique (backpropagation)[3] is used for training. Weights are updated after each backward-pass and the error (loss function) is computed

after each iteration. Once the error reduces and achieves convergence, the learning stops.

In MRSI, The time-domain complex signal of a nucleus is given by:

$$S(t) = \int p(\omega)\exp(-i\Phi)\exp(-t/T_2^*)dw. \quad (1)$$

, and the corresponding frequency-domain spectrum is given by $S(\omega)$.

Using the MLP-framework, we perform the inverse signal modeling where we have a training dataset $D = (S_i(\omega), Y_i)$, $i \in [1, N]$, where N is the total number of synthetic training spectra. $S_i(\omega)$ represents the synthetic training spectral data while Y_i represents the corresponding multi-parameter training labels. As a preliminary study, for our model, we consider the concentrations (with respect to Creatine) for the major metabolites - NAA and Choline. Therefore, for a given spectrum $S_i(\omega)$, $Y_i = [NAA_i, Cho_i]$.

A five-layered perceptron network was constructed to work as a regressor mapping the $S_i(\omega)$ to the Y_i . Each layer consisted of 300 neurons with rectified linear unit (ReLU) activation. The training data consisted of $N=1$ -million spectra with their corresponding parameters. The randomly initialized network was trained to predict the parameters by iteratively minimizing the squared-error loss between the predicted and actual parameters using gradient descent with a learning rate of 1e-3. For faster convergence, Adam optimizer with a Nesterov’s momentum of 0.9 was employed. As the data is mostly well-behaved, the ‘early stopping’ convergence check was utilized on 0.1

To check the ability of our network to predict the parameters, we use two test-sets: synthetic and real CSI. The predicted concentrations are denoted by \hat{Y}_j . The corresponding LCMoel fitted concentration labels Y_j serve as the ground-truth, $j \in [1, M]$ where M is the total number of test spectra.

Error Calculation. For our experiments, given the estimate \hat{Y}_j and the testing label Y_j , the estimate error for the parameter Y_j can be calculated as,

$$\hat{E}_j = \|\hat{Y}_j - Y_j\|/\|Y_j\| \quad (2)$$

Data. A metabolite basis set was generated by using the data provided by the ISMRM MRS Fitting Challenge 2016. An example has been shown in 1. Over 1 million spectra were simulated with variations in NAA, Cho, Cr concentrations, macro-molecular baseline, lipids, t2 values (for changes in linewidth) and signal-to-noise ratio (SNR) to account for changes in spectral quality. For testing, we acquire a standard phase-encoded 2D brain MRSI data of a healthy human volunteer on a 3T scanner using a point-resolved spin-echo localization sequence (PRESS) with voxel size = 10x10x15 mm3, TE/TR=35/1000 ms, spectral width = 2000 Hz, number of points = 400. For testing purposes, we use 96 spectra from the inner-region of the brain which serves as the region of interest. For peak alignment, ppm-cropping and signal-normalization of the training and test spectra, a pre-processing step is performed.

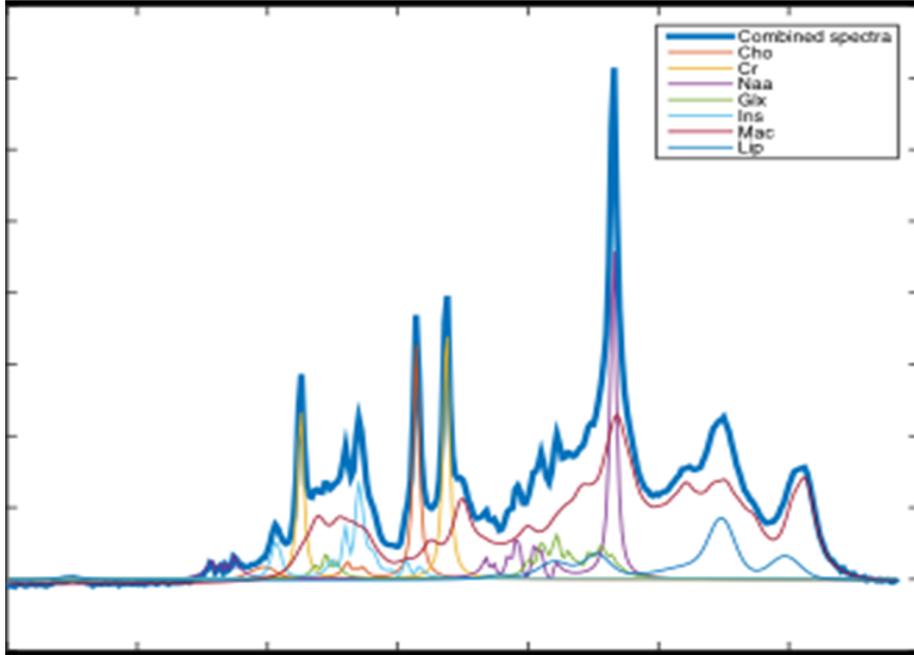


Fig. 1. An example spectra generated using the basis sets provided by the ISMRM MRS Fitting Challenge 2016. Using the same basis sets, over **1 million** spectra are generated with variations in NAA, Cho, Cr and other metabolite concentrations along with changes in macro-molecular baseline, lipids, linewidth (t_2) and SNR.

2.1 Results

As a direct comparison of both the MLP and LCModel methods, we use the synthetic test dataset to generate the error distribution shown in Fig. 2. For both NAA/Cr and Cho/Cr, the MLP shows a lower median error than the LCModel. Using the Bland-Altman method [1], we observe a strong correlation between the LCModel and RF estimates for a sample patient (Fig. 3). \hat{E}_j for the same sample patient are within the acceptable range (especially for the major metabolites such as NAA, Choline and Creatine). Fig. 4 shows the resulting concentration distribution from both the MLP and LCModel methods for both NAA/Cr and Cho/Cr. The mean relative errors over the LCModel for NAA/Cr and Cho/Cr are 0.31 and 0.12 respectively.

Speed: Training time for the synthetic data is 10 minutes using the MLP. While the LCModel takes 10 minutes for the in-vivo metabolite quantification, our proposed network, after training, takes only 10 seconds leading to a 60x improvement in speed.

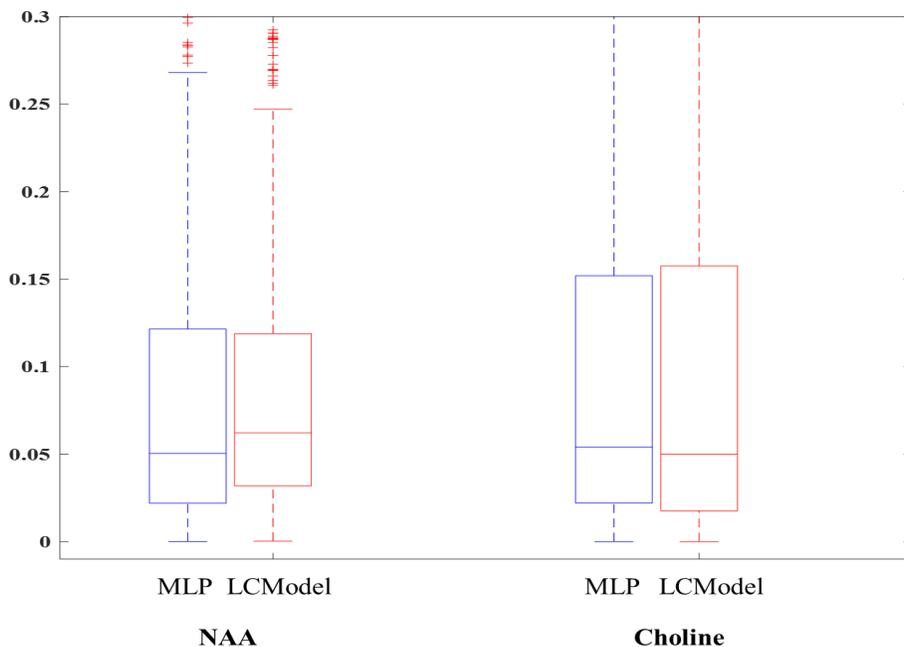


Fig. 2. Estimation error for different metabolite concentration ratios in a synthetic spectra test-set. Whiskers span the [min max] values. Median error values are represented by the red line and are as follows: NAA/Cr MLP = 0.050, LCMoDel = 0.065, Cho/Cr MLP = 0.0505, LCMoDel = 0.050.

3 Discussion and Conclusion

While the synthetic test-results gave a lower error compared to the LCMoDel, the in-vivo testing gave a slightly higher relative error. A larger training set with more training labels and a stronger network would solve this issue by providing a robust classification of real data. In our proposed method, testing and concentration estimation happens in only a few seconds and is considerably faster than the LCMoDel fitting. The deep neural-networks may be used directly, or indirectly by initializing LCMoDel fits thereby improving their results in the presence of noise and speeding up convergence.

Future work would involve using a more diverse network with layer-wise training of spectral features to improve the accuracy of parameter estimation. Once a framework has been established, further work can be done on combining these networks with global decisions about predicting spectral quality especially in the presence of artifacts.

Acknowledgements. The research leading to these results has received funding from the European Union’s H2020 Framework Programme (H2020-MSCA-ITN-2014) under grant agreement n 642685 MacSeNet.

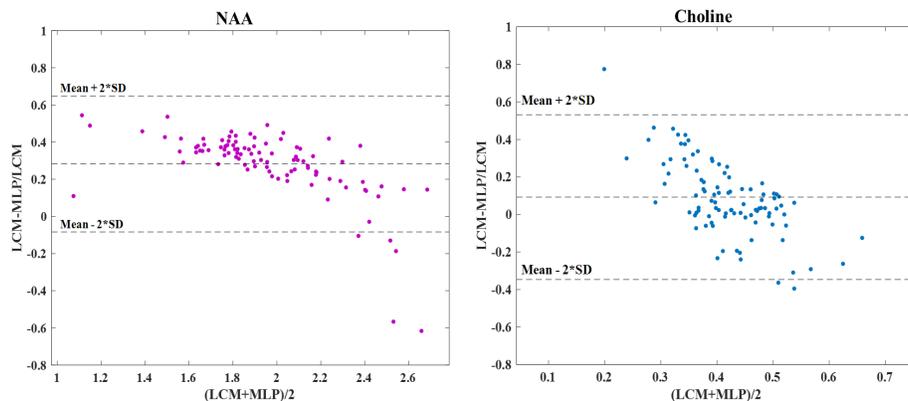


Fig. 3. Bland-Altman plots [1] representing LCModel and Multi-Layer Perceptron (MLP) estimates of spectra for the real CSI dataset. The X-Axis is the mean of the LCModel and MLP estimate, while the Y-Axis represents the relative error of the MLP estimate over the LCModel. Bland-Altman plots for: **(Left)**: NAA and **(Right)** Choline. Both plots show a good correlation with very few outliers.

References

1. Bland, J.M., Altman, D.: Statistical methods for assessing agreement between two methods of clinical measurement. *The Lancet* 327(8476), 307 – 310 (1986), <http://www.sciencedirect.com/science/article/pii/S0140673686908378>, originally published as Volume 1, Issue 8476
2. Das, D., Coello, E., Schulte, R.F., Menze, B.H.: Quantification of Metabolites in Magnetic Resonance Spectroscopic Imaging Using Machine Learning, pp. 462–470. Springer International Publishing, Cham (2017), https://doi.org/10.1007/978-3-319-66179-7_53
3. Hecht-Nielsen, R.: Theory of the backpropagation neural network. In: International 1989 Joint Conference on Neural Networks. pp. 593–605 vol.1 (1989)
4. Provencher, S.W.: Estimation of metabolite concentrations from localized in vivo proton NMR spectra. *Magnetic Resonance in Medicine* (6), 672–9
5. Schulte, R.F., Lange, T., Beck, J., Meier, D., Boesiger, P.: Improved two-dimensional J-resolved spectroscopy. *NMR in Biomedicine* 19(2), 264–270 (2006)
6. Vanhamme, L., van den Boogaart, A., Huffel, S.V.: Improved method for accurate and efficient quantification of mrs data with use of prior knowledge. *Journal of Magnetic Resonance* 129(1), 35 – 43 (1997), <http://www.sciencedirect.com/science/article/pii/S1090780797912441>
7. Widrow, B., Lehr, M.A.: 30 years of adaptive neural networks: perceptron, madaline, and backpropagation. *Proceedings of the IEEE* 78(9), 1415–1442 (Sep 1990)
8. Wilson, M., Reynolds, G., Kauppinen, R.A., Arvanitis, T.N., Peet, A.C.: A constrained leastsquares approach to the automated quantitation of in vivo 1h magnetic resonance spectroscopy data. *Magnetic Resonance in Medicine* 65(1), 1–12 (1 2011), <http://dx.doi.org/10.1002/mrm.22579>

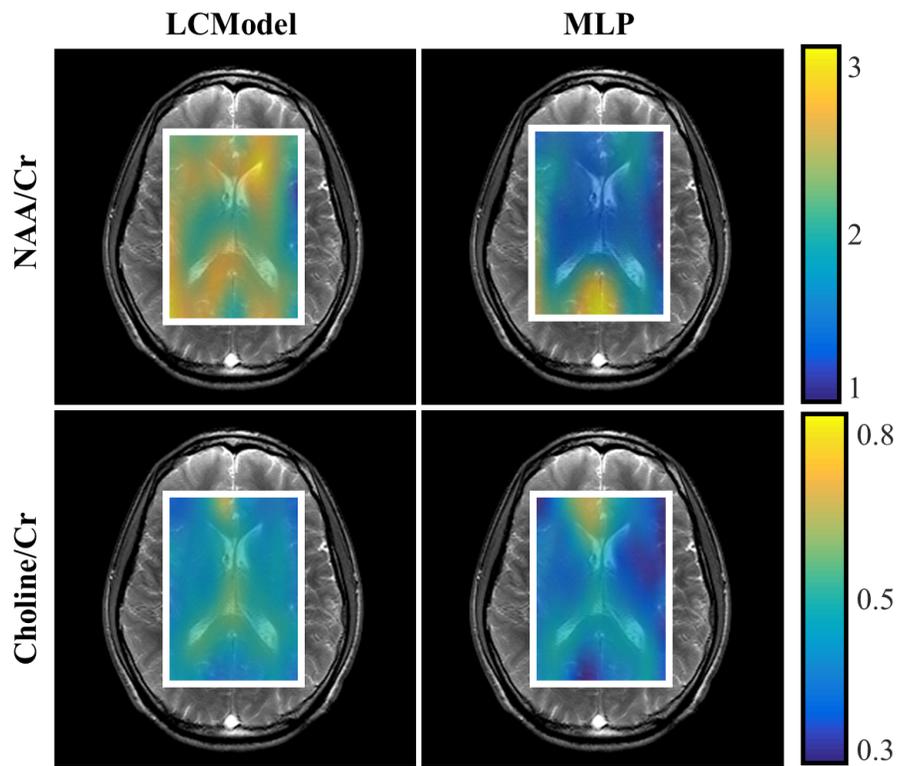


Fig. 4. NAA/Cr and Cho/Cr concentration distribution estimates from (Left)LCModel fit and (Right) multi-layer perceptron (MLP). The mean relative errors are 0.31 for NAA/Cr and 0.12 for Cho/Cr.



Bibliography

- [1] Robin A. de Graaf. *In Vivo NMR Spectroscopy: Principles and Techniques*. 2nd ed. John Wiley & Sons, 2013, 2013, p. 592.
- [2] Y. Wang and S.-j. Li. “Differentiation of Metabolic Concentrations Between Gray Matter and White Matter of Human Brain by *in vivo* ^1H Magnetic Resonance Spectroscopy.” In: ().
- [3] J. R. Alger. “Quantitative Proton Magnetic Resonance Spectroscopy and Spectroscopic Imaging of the Brain A Didactic Review.” In: 21.2 (2011), pp. 115–128.
- [4] R. Mekte, V. Mlynárik, G. Gambarota, M. Hergt, G. Krueger, and R. Gruetter. “MR spectroscopy of the human brain with enhanced signal intensity at ultrashort echo times on a clinical platform at 3T and 7T.” In: *Magnetic Resonance in Medicine* 61.6 (2009), pp. 1279–1285.
- [5] G. L. Chadzynski and U. Klose. “Chemical shift imaging without water suppression at 3 T.” In: *Magnetic Resonance Imaging* 28.5 (2010), pp. 669–675.
- [6] Z. Dong and B. Peterson. “The rapid and automatic combination of proton MRSI data using multi-channel coils without water suppression.” In: *Magnetic Resonance Imaging* 25.8 (2007), pp. 1148–1154.
- [7] S. W. Provencher. “Estimation of metabolite concentrations from localized *in vivo* proton NMR spectra.” In: *Magnetic Resonance in Medicine* 30.6 (1993), pp. 672–9.
- [8] B. M. Kelm, F. O. Kaster, A. Henning, M. A. Weber, P. Bachert, P. Boesiger, F. A. Hamprecht, and B. H. Menze. “Using spatial prior knowledge in the spectral fitting of MRS images.” In: *NMR in Biomedicine* 25.1 (2012), pp. 1–13.

- [9] D. Stefan, F. D. Cesare, A. Andrasescu, E. Popa, A. Lazariev, E. Vescovo, O. Strbak, S. Williams, Z. Starcuk, M. Cabanas, D. van Ormondt, and D. Graveron-Demilly. “Quantitation of magnetic resonance spectroscopy signals: the jMRUI software package.” In: *Measurement Science and Technology* 20.10 (2009), p. 104035.
- [10] M. Wilson, G. Reynolds, R. A. Kauppinen, T. N. Arvanitis, and A. C. Peet. “A Constrained Least-Squares Approach to the Automated Quantitation of In Vivo (1)H Magnetic Resonance Spectroscopy Data.” In: *Magnetic Resonance in Medicine* 65.1 (2011), pp. 1–12.
- [11] H. M. Nguyen, X. Peng, M. N. Do, and Z.-p. Liang. “LOW-RANK APPROXIMATIONS.” In: 3 (2011), pp. 857–860.
- [12] a. Buades and B. Coll. “A non-local algorithm for image denoising.” In: *Computer Vision and Pattern* 2 (2005), pp. 60–65.
- [13] P. Coupé, P. Yger, S. Prima, P. Hellier, and C. Kervrann. “An Optimized Blockwise Non Local Means Denoising Filter for 3D Magnetic Resonance Images An Optimized Blockwise Non Local Means Denoising Filter for 3D Magnetic Resonance Images.” In: *IEEE transactions on medical imaging* 27.4 (2008), pp. 425–441.
- [14] F. Lugauer, D. Nickel, J. Wetzl, S. A. R. Kannengiesser, A. Maier, and J. Hornegger. “Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015.” In: 9350 (2015), pp. 667–674.
- [15] S. H. Joshi, A. Marquina, S. Njau, K. L. Narr, and R. P. Woods. “DENOISING OF MR SPECTROSCOPY SIGNALS USING TOTAL VARIATION AND ITERATIVE GAUSS-SEIDEL GRADIENT UPDATES Ahmanson-Lovelace Brain Mapping Center , University of California at Los Angeles , USA 2 Department of Mathematics , University of Valencia , Spain.” In: 1 (2015), pp. 576–579.
- [16] W. Xia, E. Maneas, D. I. Nikitichev, C. A. Mosse, G. Sato dos Santos, T. Vercauteren, A. L. David, J. Deprest, S. Ourselin, P. C. Beard, and A. E. Desjardins. “Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015.” In: *MICCAI 2015, Part I, Lecture Notes in Computer Science* 9349 (2015), pp. 371–378.

-
- [17] A. Laruelo, L. Chaari, H. Batatia, S. Ken, B. Rowland, A. Laprie, and J. Y. Tourneret. “Hybrid sparse regularization for magnetic resonance spectroscopy.” In: *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBS 1* (2013), pp. 6768–6771.
- [18] B. H. Menze, B. M. Kelm, M.-A. Weber, P. Bachert, and F. A. Hamprecht. “Mimicking the human expert: Pattern recognition for an automated assessment of data quality in MR spectroscopic images.” In: *Magnetic Resonance in Medicine* 59.6 (2008), pp. 1457–1466.
- [19] D. Das, E. Coello, R. F. Schulte, and B. H. Menze. “Spatially Adaptive Spectral Denoising for MR Spectroscopic Imaging using Frequency-Phase Non-Local Means.” In: *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2016*. Ed. by S. Ourselin, L. Joskowicz, M. R. Sabuncu, G. Unal, and W. Wells. Cham: Springer International Publishing, 2016, pp. 596–604.
- [20] D. Das, E. Coello, R. F. Schulte, and B. H. Menze. “Quantification of Metabolites in Magnetic Resonance Spectroscopic Imaging Using Machine Learning.” In: *Medical Image Computing and Computer Assisted Intervention, MICCAI 2017*. Ed. by M. Descoteaux, L. Maier-Hein, A. Franz, P. Jannin, D. L. Collins, and S. Duchesne. Cham: Springer International Publishing, 2017, pp. 462–470.
- [21] D. Das, E. Coello, A. Sekuboyina, R. F. Schulte, and B. H. Menze. “Direct estimation of model parameters in MR spectroscopic imaging using deep neural networks.” In: *Proc Intl Soc Mag Reson Med(2018), Paris, France*. 2018.
- [22] D. J. MacKay. *Information theory, inference and learning algorithms*. Cambridge university press, 2003.
- [23] C. M. Bishop. *Pattern recognition and machine learning*. Springer-Verlag New York, 2006.
- [24] T. Hastie, R. Tibshirani, and J. Friedman. *The elements of statistical learning: data mining, inference and prediction*. 2nd ed. Springer, 2009.
- [25] D. Koller and N. Friedman. *Probabilistic graphical models: principles and techniques*. MIT press, 2009.

BIBLIOGRAPHY

- [26] I. Goodfellow, Y. Bengio, and A. Courville. *Deep learning*. MIT press, 2016.
- [27] T. Lange, R. F. Schulte, and P. Boesiger. “Quantitative J-resolved prostate spectroscopy using two-dimensional prior-knowledge fitting.” In: *Magnetic Resonance in Medicine* 59.5 (2008), pp. 966–972.
- [28] S. Posse, C. DeCarli, and D. Le Bihan. “Three-dimensional echo-planar MR spectroscopic imaging at short echo times in the human brain.” In: *Radiology* 192.3 (1994), pp. 733–738.
- [29] M. Lustig, D. Donoho, and J. M. Pauly. “Sparse MRI: The application of compressed sensing for rapid MR imaging.” In: *Magnetic Resonance in Medicine* 58.6 (2007), pp. 1182–1195.
- [30] S. Hu, M. Lustig, A. P. Chen, J. Crane, A. Kerr, D. A. C. Kelley, R. Hurd, J. Kurhanewicz, S. J. Nelson, J. M. Pauly, and D. B. Vigneron. “Compressed sensing for resolution enhancement of hyperpolarized ^{13}C flyback 3D-MRSI.” In: *Journal of Magnetic Resonance* 192.2 (2008), pp. 258–264.
- [31] J. K. Furuyama, N. E. Wilson, B. L. Burns, R. Nagarajan, D. J. Margolis, and M. A. Thomas. “Application of compressed sensing to multidimensional spectroscopic imaging in human prostate.” In: *Magnetic resonance in medicine : official journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine* 67.6 (2012), pp. 1499–505.
- [32] D. L. Donoho. “Compressed sensing.” In: *IEEE Transactions on Information Theory* 52.4 (2006), pp. 1289–1306.
- [33] J. Ma, Z. Sun, G. Dong, and G. Xie. “Wavelet Denoise on MRS Data Based on ICA and PCA.” In: *Ic* (2005), pp. 748–753.
- [34] R. Stoyanova, A. Kueswl, and T. R. Brown. *Application of Principal-Component Analysis for NMR Spectral Quantitation*. 1995.
- [35] F. Lugauer, D. Nickel, J. Wetzl, S. A. R. Kannengiesser, A. Maier, and J. Hornegger. “LNCS 9350 - Robust Spectral Denoising for Water-Fat Separation in Magnetic Resonance Imaging.” In: (2015).

-
- [36] Y. Liu, C. Ma, B. A. Clifford, F. Lam, C. L. Johnson, and Z. P. Liang. “Improved low-rank filtering of magnetic resonance spectroscopic imaging data corrupted by noise and B0 field inhomogeneity.” In: *IEEE Transactions on Biomedical Engineering* 63.4 (2016), pp. 841–849.
- [37] F. Lam and Z. P. Liang. “A subspace approach to high-resolution spectroscopic imaging.” In: *Magnetic Resonance in Medicine* 71.4 (2014), pp. 1349–1357.
- [38] Zheng-Hua Wu, Fan Lam, Chao Ma, and Zhi-Pei Liang. “Improved image reconstruction for subspace-based spectroscopic imaging using non-quadratic regularization.” In: *Proceedings of the 2014 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. 2014 (2014), pp. 2432–2435.
- [39] F. Lam, C. Ma, B. Clifford, C. L. Johnson, and Z. P. Liang. “High-resolution 1H-MRSI of the brain using SPICE: Data acquisition and image reconstruction.” In: *Magnetic Resonance in Medicine* 00.September (2015), pp. 1–12.
- [40] C. Ma, F. Lam, Q. Ning, C. L. Johnson, and Z. P. Liang. “High-resolution 1H-MRSI of the brain using short-TE SPICE.” In: *Magnetic Resonance in Medicine* 00 (2016).
- [41] 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.” In: *Magnetic Resonance in Medicine* 75.2 (2016), pp. 488–497.
- [42] C. Ulas, C. Preibisch, J. Sperl, T. Pyka, J. Kalpathy-Cramer, and B. H. Menze. “Accelerated Reconstruction of Perfusion-Weighted MRI Enforcing Jointly Local and Nonlocal Spatio-temporal Constraints.” In: *CoRR* abs/1708.07808 (2017).
- [43] Y. Rathi, L. Ning, O. Michailovich, H. Liao, B. Gagoski, P. E. Grant, M. E. Shenton, R. Stern, C. F. Westin, and A. Lin. “Maximum entropy estimation of glutamate and glutamine in MR spectroscopic imaging.” In: *Medical Image Computing and Computer-Assisted Intervention* 17.Pt 2 (2014), pp. 749–756.

- [44] J. V. Manjn, P. Coup, A. Buades, V. Fonov, D. Louis Collins, and M. Robles. “Non-local MRI upsampling.” In: *Medical Image Analysis* 14.6 (2010), pp. 784–792.
- [45] K. K. Bhatia, J. Caballero, A. N. Price, Y. Sun, J. V. Hajnal, and D. Rueckert. “Medical Image Computing and Computer-Assisted Intervention — MICCAI 2015.” In: 9351 (2015), pp. 510–518.
- [46] S. R. Fanello, C. Keskin, P. Kohli, S. Izadi, J. Shotton, A. Criminisi, U. Pattacini, and T. Paek. “Filter forests for learning data-dependent convolutional kernels.” In: *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (2014), pp. 1709–1716.
- [47] J. T. Bushberg. *The Essential Physics for Medical Imaging*. 2000.
- [48] C. Cortes and V. Vapnik. “Support-vector networks.” In: *Machine learning* 20.3 (1995), pp. 273–297.
- [49] L. Breiman. “Random forests.” In: *Machine learning* 45.1 (2001), pp. 5–32.
- [50] F. Rosenblatt. “The perceptron: A probabilistic model for information storage and organization in the brain.” In: *Psychological review* 65.6 (1958), p. 386.
- [51] G. Cybenko. “Approximation by superpositions of a sigmoidal function.” In: *Mathematics of Control, Signals, and Systems (MCSS)* 2.4 (1989), pp. 303–314.
- [52] A. Krizhevsky, I. Sutskever, and G. E. Hinton. “Imagenet classification with deep convolutional neural networks.” In: *Advances in neural information processing systems*. 2012, pp. 1097–1105.
- [53] S. Hochreiter and J. Schmidhuber. “Long short-term memory.” In: *Neural computation* 9.8 (1997), pp. 1735–1780.
- [54] P. Vincent, H. Larochelle, I. Lajoie, Y. Bengio, and P.-A. Manzagol. “Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion.” In: *Journal of Machine Learning Research* 11.Dec (2010), pp. 3371–3408.
- [55] K. He, X. Zhang, S. Ren, and J. Sun. “Deep residual learning for image recognition.” In: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2016, pp. 770–778.

-
- [56] B. H. Menze, A. Jakab, S. Bauer, J. Kalpathy-Cramer, K. Farahani, J. Kirby, Y. Burren, N. Porz, J. Slotboom, R. Wiest, et al. “The multimodal brain tumor image segmentation benchmark (BRATS).” In: *IEEE transactions on medical imaging* 34.10 (2015), pp. 1993–2024.
- [57] *LiTS – Liver Tumor Segmentation Challenge*. 2018. URL: <https://competitions.codalab.org/competitions/17094> (visited on 03/01/2018).
- [58] O. Maier, B. H. Menze, J. von der Gablentz, L. Häni, M. P. Heinrich, M. Liebrand, S. Winzeck, A. Basit, P. Bentley, L. Chen, et al. “ISLES 2015-A public evaluation benchmark for ischemic stroke lesion segmentation from multispectral MRI.” In: *Medical image analysis* 35 (2017), pp. 250–269.
- [59] B. H. Menze, B. M. Kelm, M.-A. Weber, P. Bachert, and F. A. Hamprecht. “Mimicking the human expert: Pattern recognition for an automated assessment of data quality in MR spectroscopic images.” In: *Magnetic Resonance in Medicine* 59.6 (2008), pp. 1457–1466.
- [60] N. Pedrosa de Barros, R. McKinley, R. Wiest, and J. Slotboom. “Improving labeling efficiency in automatic quality control of MRSI data.” In: *Magnetic Resonance in Medicine* (2017), p. 1.
- [61] D. L. Collins, a. P. Zijdenbos, V. Kollokian, J. G. Sled, N. J. Kabani, C. J. Holmes, and a. C. Evans. “Design and construction of a realistic digital brain phantom.” In: *IEEE transactions on medical imaging* 17.3 (1998), pp. 463–468.
- [62] *Vespa Project (Versatile Simulation, Pulses and Analysis)*.
- [63] P. J. W. Pouwels and T. Frahm. “Regional Metabolite Concentrations in Human Brain as Determined by Quantitative Localized Proton MRS.” In: (), pp. 53–60.
- [64] P. A. Bottomley. “Spatial Localization in NMR Spectroscopy in Vivo.” In: *Annals of the New York Academy of Sciences* 508 (1987), pp. 333–348.
- [65] A. Haase, J. Frahm, W. Hänicke, and D. Matthaei. “ ^1H NMR chemical shift selective (CHESS) imaging.” In: *Physics in medicine and biology* 30.4 (1985), pp. 341–344.

BIBLIOGRAPHY

- [66] L. Breiman. “Random Forests.” In: *Machine Learning* 45.1 (2001), pp. 5–32.
- [67] A. Fuchs, P. Boesiger, R. F. Schulte, and A. Henning. “ProFit revisited.” In: *Magnetic Resonance in Medicine* 71.2 (2014), pp. 458–468.
- [68] ”. M. Bland and D. Altman”. “”Statistical Methods for Assessing Agreement between Two Methods of Clinical Measurement”.” In: *”The Lancet” ”327” .”8476” (”1986”)*. ”Originally published as Volume 1, Issue 8476” , ”307–310” .
- [69] B. H. Menze, A. Jakab, S. Bauer, J. Kalpathy-Cramer, K. Farahani, J. Kirby, Y. Burren, N. Porz, J. Slotboom, R. Wiest, et al. “The multimodal brain tumor image segmentation benchmark (BRATS).” In: *IEEE transactions on medical imaging* 34.10 (2015), pp. 1993–2024.
- [70] F. Knoll, J. Zbontar, A. Sriram, M. J. Muckley, M. Bruno, A. Defazio, M. Parente, K. J. Geras, J. Katsnelson, H. Chandarana, Z. Zhang, M. Drozdalv, A. Romero, M. Rabbat, P. Vincent, J. Pinkerton, D. Wang, N. Yakubova, E. Owens, C. L. Zitnick, M. P. Recht, D. K. Sodickson, and Y. W. Lui. “fastMRI: A Publicly Available Raw k-Space and DICOM Dataset of Knee Images for Accelerated MR Image Reconstruction Using Machine Learning.” In: *Radiology: Artificial Intelligence* 2.1 (2020), e190007.
- [71] I. Marshall, M. J. Thrippleton, M. E. Bastin, D. Mollison, D. A. Dickie, F. M. Chappell, S. I. K. Semple, A. Cooper, S. Pavitt, G. Giovannoni, C. A. M. G. Wheeler-Kingshott, B. S. Solanky, C. J. Weir, N. Stallard, C. Hawkins, B. Sharrack, J. Chataway, P. Connick, S. Chandran, and for the MS-SMART Trialists. “Characterisation of tissue-type metabolic content in secondary progressive multiple sclerosis: a magnetic resonance spectroscopic imaging study.” In: *Journal of Neurology* 265.8 (Aug. 1, 2018), pp. 1795–1802.