

# MRSNet: Metabolite Quantification from Edited Magnetic Resonance Spectra With Convolutional Neural Networks

M. Chandler, C. Jenkins, S. M. Shermer, and F. C. Langbein

**Abstract**—Quantification of metabolites from magnetic resonance spectra (MRS) has many applications in medicine and psychology, but remains a challenging task despite considerable research efforts. For example, the neurotransmitter  $\gamma$ -aminobutyric acid (GABA), present in very low concentration *in vivo*, regulates inhibitory neurotransmission in the brain and is involved in several processes outside the brain. Reliable quantification is required to determine its role in various physiological and pathological conditions. We present a novel approach to quantification of metabolites from MRS with convolutional neural networks — MRSNet. MRSNet is trained to perform the multi-class regression problem of identifying relative metabolite concentrations from given input spectra, focusing specifically on the quantification of GABA, which is particularly difficult to resolve. Typically it can only be detected at all using special editing acquisition sequences such as MEGA-PRESS. A large range of network structures, data representations and automatic processing methods are investigated. Results are benchmarked using experimental datasets from test objects of known composition and compared to state-of-the-art quantification methods: LCMoDel, jMRUI (AQUES, QUEST), TARQUIN, VeSPA and Gannet. The results show that the overall accuracy and precision of metabolite quantification is improved using convolutional neural networks.

**Index Terms**—Convolutional Neural Networks, Multi-class Regression, Magnetic Resonance Spectroscopy, Metabolite Quantification, GABA, MEGA-PRESS.



## 1 INTRODUCTION

MAGNETIC resonance spectroscopy (MRS) is a rapidly developing, highly versatile tool, enabling us to acquire information about biochemical processes happening *in vivo* and to identify biochemical changes associated with various cancers and neurological disorders among others. The technique is non-invasive, painless, involves no ionizing radiation or radioactive tracers and can be performed on conventional clinical MRI scanners available in many hospitals. MRS has been responsible for advances in detecting, classifying and modelling malignant tissue in the brain [1], [2], [3], prostate [4], [5], [6] and breast [7], [8], [9]. It also has a role in modern psychology and neuroscience, where it provides a useful window into normal brain function, as well as abnormalities associated with conditions such as anxiety, schizophrenia, depression and other mood disorders [10], [11], [12], [13], [14]. Many of these studies have focused on the detection of  $\gamma$ -aminobutyric acid (GABA), the primary inhibitory neurotransmitter [15], as a key biomarker. Its quantification is especially difficult as it is present in low concentration *in vivo* and its characteristic MRS features overlap with those of much more abundant metabolites, such as N-Acetylaspartic acid (NAA) and Creatine (Cr), obscuring its signature in MR spectra. This had lead to the development of new spectroscopic techniques, such as

edited spectroscopy, that attempt to selectively edit out certain features to make others observable. In this work, we focus on quantification of GABA with MEGA-PRESS, one of the most commonly used edited spectroscopy sequences.

An MR spectrum is a graph of the magnetic resonance signal  $S(\nu)$  as a function of frequency  $\nu$ . MR spectra of different molecules have different features such as a set of characteristic peaks at distinct frequencies. While individual MR spectra are highly molecule-specific, identifying and quantifying the concentration of molecules from MR spectra of mixtures containing many molecules is a very challenging task as spectral features of different molecules overlap, signals from molecules present in low concentration may be barely distinguishable from the noise floor, resulting in low signal-to-noise ratios (SNR), and hardware or software calibration issues may distort the spectra obtained. In essence, quantification decomposes spectra into features associated with individual metabolites to infer the chemical composition of a sample or tissue voxel. There are many ways to tackle quantification, such as basis set fitting, singular value decomposition, principal component analysis or peak integration, focusing on either time or frequency domain analysis. In addition, spectral processing methods are employed for phase correction, removal of water or macro-molecule signals and baseline correction. Many of these require manual calibration and spectral processing may lead to distortion of spectra, thereby falling short of achieving accurate and precise quantification with minimal input from a human expert. [16] is an excellent review of the state-of-the-art.

Basis set methods achieve decomposition by representing a spectrum as a linear combination of spectra from a

- M. Chandler and F. C. Langbein are with the School of Computer Science and Informatics, Cardiff University, Cardiff, UK.  
E-mail: hello@ma.ax, frank@langbein.org
- C. Jenkins is with CÚBRIC at Cardiff University, Cardiff, UK.  
E-mail: JenkinsC23@cardiff.ac.uk
- S. M. Shermer is with the Department of Physics at Swansea University, Swansea, UK.  
E-mail: lw1660@gmail.com

basis set of individual metabolite signals and a residual signal. State-of-the-art methods generally employ a variant of least-squares fitting to match a linear combination to the spectrum. Creating basis sets itself, whether based on simulation or measurements, is not a trivial task. Further complicating basis set selection is that basis sets must usually be specifically created for the pulse sequence used to acquire the spectra, as the spectra depend on the pulse sequence. This is problematic as there are many pulse sequences and implementations on different systems can vary considerably [17]. Deriving basis sets from experimental data is also extremely time-consuming and the resulting basis sets may suffer from the same problems as the spectra to be quantified. Alternatively, using simulations for basis set generation relies on models, the accuracy of which has been questioned [18]. Quantification, therefore, can be an extremely time-consuming task, where time must be spent acquiring or simulating basis sets as well as adjusting fitting and processing algorithms and parameters. This is where machine learning has an opportunity to improve quantification by providing a complete, unbiased solution in less time. Machine learning has already proved to be a useful tool in MR imaging [19], [20], MR spectroscopy [21], [22], [23], [24], classification of Raman [25] and Faraday spectra [26], and near-infrared (NIR) spectroscopy calibration [27].

We present MRSNet, a multi-class regression convolutional neural network (CNN) that can be trained quickly on a relatively small number of samples (5,000). We review MRS and associated quantification methods in Section 2. The data and method for the CNNs are discussed in Sections 3 and 4 respectively. We focus on frequency domain analysis by training the CNN on spectra to learn metabolite parameters, returning expected relative concentrations. Due to the time required to experimentally acquire a large number of spectra, the network is trained on a range of simulated data (see Section 5). As the results are heavily dependent on the accuracy of the basis set and the intra-basis set scaling, multiple basis sets from different sources are compared in Section 5.3. Furthermore, a range of network structures and different data representations of the frequency domain signals (real part, imaginary part and magnitude) are explored in Section 5.2. The performance of MRSNet and other state-of-the-art methods is benchmarked using experimental spectra obtained from test objects (phantoms) of known composition (see Section 6). The results show that on average CNNs are more accurate and precise. The experimental data used for the evaluation is available at [28] and the code for the networks and spectral simulation is available at [29].

## 2 BACKGROUND

We briefly review the acquisition of MR spectra and state-of-the-art quantification methods.

### 2.1 Magnetic Resonance Spectroscopy

MRS utilizes the fact that elementary particles such as protons, neutrons and electrons have a quantum-mechanical degree of freedom called spin. As a result, atomic nuclei with an odd number of protons or neutrons such as hydrogen  $^1\text{H}$  have a net nuclear spin.  $^1\text{H}$ , the most abundant

element in biological organisms [30], has two spin states. When an external magnetic field  $B_0$  is applied, an energy gap of  $h\nu$  develops between these states, where  $h$  is the Planck constant and  $\nu$  is the Larmor frequency, which is proportional to  $B_0$ . In a typical clinical MR scanner the Larmor frequency of  $^1\text{H}$  is in the radiofrequency (RF) range, e.g., 127 MHz at 3 T. As the lower energy state is energetically preferred, a slight excess of nuclei are found in this state, leading to the development of a magnetic moment in the sample. This magnetic moment is almost undetectable, but when a suitable electromagnetic pulse at the Larmor frequency is applied, the spins can be rotated into a plane transverse to the external magnetic field. As the spins precess in this transverse plane they emit an MR signal while slowly returning to their equilibrium state, which can be detected by sensitive RF receive coils.

This time-domain signal is complex, as it has both magnitude and phase, and is Fourier-transformed to obtain the MR spectrum. The reason for this is that the features in the frequency-domain representation are directly related to the structure of the molecule the signal originates from. Although the Larmor frequency of an isolated proton depends only on the external  $B_0$  field, protons in molecules have slightly shifted resonance frequencies due to chemical shielding effects, known as chemical shift. Chemical shifts can be reported in Hz but it is customary to report chemical shifts in ppm (parts-per-million) with respect to the spectrometer frequency  $\nu_0$ . For a spectrometer calibrated to 127 MHz, a chemical shift of 1 ppm corresponds to 127 Hz. Chemical shifts accessible by MR spectroscopy are usually in the range of 0 ppm to 8 ppm and can be affected by pH and temperature of the sample. In addition to chemical shift, MR spectra are also influenced by  $J$ -couplings between nearby spins. This leads to splitting of the energy levels of the coupled spin system, which manifests itself as splitting of individual peaks into doublets, triplets, quadruplets, etc., depending on the number of spins that are simultaneously coupled. In a perfectly homogeneous field and in the absence of any relaxation the peaks would be  $\delta$ -function spikes. In practice,  $T_1$  and  $T_2$  relaxation of the precessing spins leads to line broadening and Lorentzian or Gaussian spectral lineshapes. This intrinsic line broadening is further increased by local variations in the  $B_0$  field, resulting in a distribution of frequencies over an excited volume. A homogeneous  $B_0$  field correlates with a narrower linewidth, sharper peaks and, importantly, reduced feature overlap between metabolite spectra.

In practice, MRS pulse sequences are more complicated than a single excitation pulse followed by immediate readout. Additional pulses must be applied to spatially localize the signal (Point RESolved Spectroscopy, PRESS) [31] and suppress unwanted signals such as the background water signal, e.g., by applying a CHEMical Shift-Selective (CHESS) pulse to saturate the water signal, followed by gradient dephasing [32]. Moreover, timing parameters such as  $T_E$  (echo time; time between RF pulse and echo readout) and  $T_R$  (repetition time; time between corresponding points in a repeating series of pulses and echoes) are adjusted to reduce residual magnetization between scans and macromolecular contributions, among other reasons. All of these choices affect the spectra obtained.




Fig. 1: Simulated MEGA-PRESS spectra for mixture of 15 mmol/l NAA, 8 mmol/l Cr and 3 mmol/l GABA (left) and for GABA only, at much lower signal strength (right).

For edited spectroscopy, such as the MEGA-PRESS (MEsher-GARwood PRESS) sequence [33], additional editing pulses are applied between refocusing pulses to selectively enhance or suppress certain features. The standard MEGA-PRESS sequence produces two MR spectra, an edit-off and edit-on spectrum, from which a third difference spectrum is calculated. Specifically for GABA quantification, the editing pulses are applied at 1.9 ppm during the edit-on acquisition and 7.4 ppm during the edit-off acquisition. The aim of these editing pulses is the elimination of the Cr signal in the difference spectrum at 3 ppm. As shown in Fig. 1, the small GABA signal at 3 ppm is completely dominated by the much larger Cr signal. In the edit-off spectrum, the outer peaks of the GABA triplet at 3 ppm are inverted compared to the edit-on spectrum, while the Cr signal is unaffected. In the difference spectrum, this ensures that the Cr signal at 3 ppm is eliminated and only the side peaks of the 3 ppm GABA triplet remain, amplified by a factor of 2, making GABA detectable. The 2 ppm NAA peak is absent from the edit-on spectrum as the 1.9 ppm editing pulse also eliminates its corresponding magnetic moments.

## 2.2 Quantification of MR Spectra

Quantification of MR spectra aims to infer the relative contributions (concentrations) of different metabolites, accurately and precisely, even in the presence of environmental noise and low signal-to-noise ratios. While some metabolites are quite abundant and relatively easy to quantify, others of great importance, such as the inhibitory neurotransmitter GABA and the excitatory neurotransmitters Glutamine (Gln) and Glutamate (Glu) [34], are considerably more challenging. This is due to low relative concentration *in vivo* and overlap of their spectral features with those of more abundant metabolites [35]. For Glu/Gln the difficulty lies further in resolving their spectra, which are extremely similar at the field strength at which most clinical MRI scanners operate (1.5 T or 3 T). The primary goal of this paper is the quantification of these difficult-to-quantify metabolites. As absolute quantification requires reliable reference spectra for calibration, which are generally not available for *in vivo* spectra, metabolite concentrations are typically reported as ratios, such as Cr/GABA, NAA/GABA and we also focus on relative rather than absolute quantification.




Fig. 2: Simplified quantification example for spectrum containing Cr (blue), GABA (orange), Glu (green), Gln (red) and NAA (purple): simulated combined spectrum with noise (left), example fit (middle), quantification results (right).

### 2.2.1 Traditional Approaches

Various quantification methods have been developed, ranging from simple peak integration methods to fitting either the time or frequency domain signals with a linear combination of basis signals plus a residual in absolute or relative values as shown in Fig. 2. State-of-the-art methods typically focus on frequency domain analysis, due to difficulties interpolating time domain signals with varying acquisition parameters. The majority implement a variant of the non-linear least-squares (NLLS) algorithm to fit a set of parameters to a model. Methods that employ basis sets usually come bundled with a simulator to generate them. We briefly describe state-of-the-art software for quantification and simulation. For an excellent review see [36].

LCModel [37] is widely considered to be the gold standard of spectral quantification and uses NLLS to decompose spectra in the frequency domain. Basis sets are provided by the creator on request or can be obtained from suggested external sources. LCModel is only able to directly simulate macro-molecule contributions and not complete basis sets.

In contrast, TARQUIN [38], VeSPA [39] and jMRUI [40] are able to simulate their own basis sets for quantification. TARQUIN and VeSPA both perform NLLS fitting of frequency domain signals, while jMRUI has a range of algorithms available, including AMARES and AQSES for frequency domain analysis and QUEST for the time domain.

Gannet [41] focuses on peak identification and fitting specifically for edited MRS spectra, such as MEGA-PRESS, and does not use a basis set or simulator.

Despite the advances made in MR spectra quantification, it has been shown that there are significant errors in quantification between the different methods [36], [42], [43]. Recent attempts to compare and validate these tools for spectra obtained from carefully calibrated phantoms have shown very poor accuracy and large discrepancies between the concentrations reported by different tools [44], [45]. These studies highlight the need for more reliable methods for quantification of MR spectra for MRS to become a reliable tool for clinical use, especially for difficult-to-quantify metabolites such as GABA.

Other methods utilising Bayesian inference have been developed for the analysis of time-domain signals in high-field NMR [46]. These leverage prior knowledge via probability distributions and estimate ‘nuisance parameters’ such as phase shifts, decay constants and noise parameters’ such as performing Fourier based NMR methods. Due to much higher spectral resolution and the focus on molecular struc-

ture elucidation, the issues and applications of high-field NMR are different and the methods have not been applied to MRS quantification *in vivo*.

### 2.2.2 Machine Learning Approaches

Machine learning for MRS quantification has only recently begun to be explored. Das *et al.* [23] used random forests trained on 1 million simulated spectra and tested on 287 human subjects. More recently Lee *et al.* [22] trained CNNs on 40,000 spectra to quantify a range of simulated and *in vivo* data. Both studies are focused on quantification with the single acquisition PRESS sequence, which generally precludes identification and quantification of difficult-to-detect metabolites such as GABA. Furthermore, the performance of the machine learning algorithms on *in vivo* spectra was assessed by comparing the quantification results with LCModel. Considering the large discrepancies in metabolite concentrations reported by different tools when provided with *identical* input spectra, and the overall *lack of accuracy* of the quantification results obtained with existing tools when tested against a wide range of spectra from calibrated phantoms, good agreement with LCModel results does not provide a strong validation of the method [44], [45].

Here we focus on the quantification of difficult-to-quantify metabolites in MEGA-PRESS spectra. The networks are trained on a comparatively small number of simulated spectra (5,000) and the results are validated independently using calibrated phantom spectra. A range of network structures, spectra representations, processing methods and the impact of the basis set choice are explored.

## 3 DATA

Spectral datasets for quantifying mixtures of NAA, Gln, Glu, GABA and Cr are used due to their importance for MRS applications and the difficulty they pose for accurate quantification. These metabolites are in principle detectable by MEGA-PRESS. The networks are trained and validated using simulated data generated by combining simulated spectral basis sets into spectra based on the concentration of individual metabolites. Simulated spectral basis sets were deemed more efficient than deriving basis sets experimentally from calibrated phantoms. Using experimental data to generate basis sets is also problematic due to the difficulty of obtaining high-quality spectra [47]. To assess the performance of the networks, experimental data from carefully calibrated phantoms are used as benchmark.

### 3.1 Training and Validation Dataset Generation

A basis set contains the normalised characteristic signals in the frequency domain for each metabolite, which are combined to create mixed spectra.

Basic simulators for basis sets often assume the molecule is prepared in a well-defined initial state such as the ground state. They approximate the pulse sequence by a series of instantaneous unitary operations acting on the state at certain times, followed by readout of the MR signal, to simulate the full MRS pulse sequence. More sophisticated simulators perform time-resolved calculations for finite-duration pulses and include the effects of relaxation or

field inhomogeneities. However, the choices of pulse shapes and pulse timings in vendor-specific implementations of the MEGA-PRESS sequence differ [17] and are often not known precisely. Therefore, quite often educated guesses must be made. Similarly, other parameters, such as relaxation parameters or field inhomogeneities, are typically not known precisely. Moreover, even metabolite models, including chemical shift and  $J$ -coupling parameters, still have uncertainties [18]. A major source of these parameters for common metabolites in MRS is the landmark paper by Govindaraju *et al.* [35], but alternative models have been suggested, e.g., for GABA by Near *et al.* [48] and Kaiser *et al.* [49].

Hence, we use and test different basis sets (see Section 6.2) from state-of-the-art simulation and quantification software: FID-A [50], PyGamma [51] and LCModel [37]. The MEGA-PRESS basis set for LCModel is generated by the MRS Lab at Purcell Health Sciences [52]. For PyGamma and FID-A, simulations are performed in house by adapting the MEGA-PRESS pulse sequence code provided. All simulations are aimed at the Siemens WIP MEGA-PRESS implementation, which was used to acquire the experimental datasets (see Section 3.2). These basis sets were also used in a comparative study of state-of-the-art quantification methods [44], [45] and the comparison with such methods in Section 6.1.

For each basis set, the training and validation datasets are generated by taking linear combinations of the basis elements with each metabolite having a scaling factor in  $[0, 1]$ , corresponding to a relative concentration. The scaling factors are sampled using a low discrepancy (quasi-Monte Carlo) Sobol sequence [53], which provides good uniform coverage of possible states with a low number of data points. Time-domain noise is added from a normal distribution ( $\mu = 0$  and  $\sigma$  randomly chosen in the range  $[0, 0.25]$ ) to 50% of the dataset to improve simulation accuracy and network robustness. This noise model was chosen as it closely resembles what is seen in experimental spectra, by characterising the noise profile from spectral areas that do not contain a metabolic signal for 4, 160 phantom spectra.

### 3.2 Experimental Benchmark Datasets

The performance of the network is evaluated and compared with other state-of-the-art methods using a spectral benchmark dataset created by scanning calibrated phantoms of the metabolites studied here [28]. The MEGA-PRESS spectra were acquired on a 3 T Siemens Magnetom Skyra system at Swansea University using the Siemens WIP MEGA-PRESS implementation with  $T_E = 68$  ms and  $T_R = 2000$  ms with 160 averages per spectrum. Four datasets for phantoms of known composition, E1, E3 & E4a and E4b, were used. E1, E3 and E4a were acquired with an acquisition bandwidth of 1250 Hz, while E4b was acquired at 2000 Hz (for the same phantoms as E4a). The published experimental dataset also contains a non-pH calibrated dataset E2, which was excluded as it was deemed not representative of the *in vivo* environment. E4c and E4d, also contained in the dataset, are repeat runs of E4a and E4b and we obtained similar results for them, so they are not further discussed here. Phantom composition is displayed in Table 1 and Fig. 3; for details of phantom preparation see [44], [45].

TABLE 1: Benchmark phantom composition: concentrations are in mM = mmol/l.

Series	Medium	# of Spectra	NAA	Cr	Glu	Gln	GABA
E1	Water	13	15.0	0/8.0	0.0	0.0	0.0, 0.5, 1.0, 1.5, 2.0, 2.5, 3.0, 4.0, 6.0, 8.0, 10.0, 11.6
E3	Water	15	15.0	8.0	12.0	3.0	0.0, 1.0, 2.0, 3.0, 4.0, 5.0, 6.0, 7.0, 8.0, 9.0, 10.0, 11.0, 12.0, 13.0, 14.0
E4	Gel	8	15.0	8.0	12.0	3.0	0.0, 1.5, 3.0, 5.0, 7.0, 10.0




Fig. 3: Metabolite concentrations relative to NAA, for datasets E1, E3 and E4 for NAA (purple: 1), Cr (blue: 0/0.5), GABA (orange: variable), Gln (red: 0.8), Glu (green: 0.2).

## 4 METHOD

We present the network structure and training procedure for MRSNet, our deep MRS quantification network, minimal data pre-processing and performance evaluation measures.

### 4.1 Network Structure

MRSNet is a 2D CNN. Input data to the network is a  $N \times 2048$  matrix, where each row contains one component of a single spectrum. The maximum number of input rows is  $N = 9$  for the real and imaginary part and magnitude of up to three spectra (edit-off, edit-on and difference). Spectra are zero-padded and trimmed to obtain a length of 2048 in the range from 4.5 ppm to 1.5 ppm, i.e. each column is a single frequency bin 0.0014 ppm wide, providing ample resolution for small features. The majority of convolutions are  $1 \times N$ , spanning across the frequency bins to align with the separate rows of the input spectra. This network structure is loosely based on the CNN VGG architecture [54] and experimentally derived. Multiple layer 1D and 2D networks and single-layer 3D networks were explored by varying the arrangement of acquisitions and data type representations. However, they generally have poor performance compared to the suggested network structure.

The architecture of the network is designed to train identification of spectral features in individual rows (component of a single spectrum) early on, to be combined further down the network. These reduction layers are implemented in the middle of the network and repeated to combine the rows until the output tensor has one row. Using convolutions spanning multiple rows of the input early on lead to very poor results, whereas combining the convoluted rows near the middle of the network provided the best results.

The complete network architecture is shown in Table 2 for the ‘small’ variant without pooling. For every variant of the network, the layer class *reduction1* is repeated until the output is a tensor with one row. For input spectra with nine rows, the layer is repeated five times, whereas for an input of a tensor with one row the layer has a convolutional kernel sized  $3 \times 1$  and is repeated once. This is where the

TABLE 2: Small network structure without pooling. Elipsis indicate repeated layers of the same class. Keys: conv: convolution; BN: batch normalisation; DO: drop out; FC: fully connected; ReLU: rectified linear unit.

Layer class	Description
input	$[1 - 9] \times 2048$ spectra
conv1	$256 \times 1 \times 7$ conv, $1 \times 2$ stride, ReLU, BN, DO= 0.4
conv2	$256 \times 1 \times 5$ conv, $1 \times 2$ stride, ReLU, DO= 0.4
reduction1	$256 \times [1 - 3] \times [3]$ conv, ReLU, DO= 0.25
reduction1	...
conv3	$256 \times 1 \times 3$ conv, padding=same, ReLU, DO= 0.25
reduction2	$256 \times 1 \times 3$ conv, $1 \times 3$ stride, ReLU, DO= 0.25
conv3	...
reduction2	...
conv4	$512 \times 1 \times 3$ conv, padding=same, ReLU, DO= 0.25
reduction3	$512 \times 1 \times 3$ conv, $1 \times 3$ stride, ReLU, DO= 0.25
conv4	...
reduction3	...
dense1	1024 FC Layer
output	5 FC Layer, softmax

TABLE 3: Network layer class substitutions for ‘medium’ and ‘large’ networks, differences are in bold.

Network	Layer class	Description
Medium	conv1	$256 \times 1 \times \mathbf{9}$ conv, $1 \times 2$ stride, ReLU, BN, DO= 0.4
	conv2	$256 \times 1 \times \mathbf{7}$ conv, $1 \times 2$ stride, ReLU, DO= 0.4
	reduction1	$256 \times [1 - 3] \times \mathbf{5}$ conv, ReLU, DO= 0.25
Large	conv1	$256 \times 1 \times \mathbf{16}$ conv, $1 \times 2$ stride, ReLU, BN, DO= 0.4
	conv2	$256 \times 1 \times \mathbf{8}$ conv, $1 \times 2$ stride, ReLU, DO= 0.4
	reduction1	$256 \times [1 - 3] \times \mathbf{7}$ conv, ReLU, DO= 0.25

information contained in the different rows (spectra) of the input tensor are combined. A variety of substitutions are made to test different ‘medium’ and ‘large’ networks and pooling variants as shown in Tables 3 and 4.

The network is trained using the ADAM weight update [55] with an experimentally derived learning rate of  $10^{-4}$  and suggested beta values of  $\beta_1 = 0.9$  and  $\beta_2 = 0.999$  respectively. Mean squared error (MSE) is used as the loss function over mean absolute percentage error (MAPE) or mean absolute error (MAE) due to the higher overall performance seen experimentally. The training was performed on an NVIDIA Titan X, using Python 2.7.15, Keras 2.2.4 [56] with Tensorflow 1.9.0 and CUDA 9.0.176.

### 4.2 Pre-processing

All spectra are  $B_0$  corrected w.r.t. the 2 ppm NAA singlet peak due to NAA’s low sensitivity to temperature of 0.01 ppm/ $^{\circ}\text{C}$  [57]. The experimental spectra are filtered

TABLE 4: Network substitutions for pooling variants where strided convolutions are replaced by a MaxPool layer following a convolution layer.

Network	Layer class	Description
All (pooling)	reduction2	256×1×3 conv, ReLU, DO= 0.25 1×3 MaxPool
	reduction3	512×1×3 conv, ReLU, DO= 0.25 1×3 MaxPool

using a first-order Butterworth filter to reduce noise and with no phase correction done for simplicity, as it can often require human interaction [58]. Experimental and simulated spectra are mean-centred and normalised so that the largest peak has amplitude  $\pm 1$  across acquisitions with  $x' = x/\max(|x|)$ . Different acquisition bandwidths are dealt with by zero-filling the time domain signal to achieve a spectral resolution of 2048 points in the selected ppm range.

### 4.3 Performance Evaluation

The error  $\epsilon$  is calculated as the mean of the absolute differences of the actual  $a_{i,j}$  and the predicted  $p_{i,j}$  relative concentration for every label (metabolite)  $j = 1, \dots, L$ , for every prediction (spectrum)  $i = 1, \dots, N$ :

$$\epsilon = \frac{1}{NL} \sum_{i=1}^N \sum_{j=1}^L |a_{i,j} - p_{i,j}|. \quad (1)$$

The standard deviation  $\sigma$  is calculated in the usual way,  $\sigma^2 = \sum_{i=1}^N \sum_{j=1}^L (p_{i,j} - \epsilon)^2 / (NL)$ . This provides a good indication of overall network performance but is insensitive to low concentration metabolites. To counter this, an in-depth regression and MAPE analysis is performed for the best network on a per-metabolite basis in Section 6.1.

## 5 EXPERIMENTS

We investigate a range of network structures, spectra, representations and basis sets to find a favourable combination. Network structure and batch sizes are investigated first, followed by the effect of different representations of the spectra and the choice of basis set. The LCModel basis set is used for the first two experiments as the analysis program has been shown to perform well by the comparative study in [44], [45] and it performs best compared to other state-of-the-art quantification methods (see Table 9), suggesting a good fit with experimental data. Training and validation are performed using simulated spectra as described in Section 3.1. Initial experiments in Sections 5.1 and 5.2 are trained for 100 epochs, with 4,000 training, 1,000 validation samples. For the final experiment in Section 5.3, more compute time is dedicated to refining the networks over 200 epochs with 5,000 training samples. A light early stopping criteria is utilised throughout to prevent overfitting, with a minimum loss decrease of  $10^{-12}$ , a patience (number of epochs without improvement) of 15 and ‘restore best weights’ enabled. Network performance is benchmarked using experimental datasets from samples with known composition rather than performance on validation as described in Section 3.2.

### 5.1 Network Structure Investigation

An initial investigation is performed to explore the effect of minibatch sizes, convolutional kernel widths and pooling vs. stridden convolutions on performance with the benchmark datasets. Three variants of the network are tested (small, medium and large) over a range of minibatch sizes (64, 32, 16) for two major variants of each network, either using stridden convolutions or max-pooling. These two variants are explored as convolutions with strides have been shown to be advantageous over max-pooling methods [59]. Networks are trained using all three spectra (edit-off, edit-on and difference) and all three data representations (real, imaginary and magnitude/absolute value) stacked to form a  $9 \times 2048$  input tensor for each sample. The results in Tables 5 and 6 show the general trend that using convolutions with strides obtains a substantially higher performance over max-pooling. Networks with smaller convolutional kernels are substantially quicker than their larger counterparts as they are less computationally expensive to calculate. The ‘small’ network, using stridden convolutions and a minibatch size of 64 achieves a good balance between performance and training time on the benchmark datasets, making it the chosen network architecture for the following experiments.

### 5.2 Effect of Different Representations of Spectra

Using MEGA-PRESS data provides a unique opportunity to explore how training networks on different combinations of acquired spectra for one scan (edit-off, edit-on and difference) affects performance. Additionally, three data types have been chosen to represent the spectrum by taking the real, imaginary or magnitude component of the frequency domain signal. These components are explored by using the same generated dataset for training and validation but considering different combinations of spectra and data types.

Results in Table 7 show a clear performance advantage of utilising the magnitude representation for all acquisition types. This is expected, as there is a level of phase uncertainty from the experimental signal leading to a potential disagreement with the real and phase spectra from the benchmark set when compared to the basis set. However, using the magnitude spectra does increase the linewidth when compared to only the real spectrum; which is typically preferred for quantification for this reason. Despite this, the networks prefer the more predictable magnitude spectra, at the cost of dealing with the increased linewidth.

The best combination of acquisitions is the edit-off and difference, as these provide the maximal amount of information with no repetition of data. The reduced performance for the difference only acquisition can be attributed to the Cr spectra, wherein the LCModel basis set that was used to train the network assumes perfect editing, where the normal and inverted spectrum in the edit-on and edit-off acquisition match amplitude. As such, there is no residual Cr signal in the difference spectrum (which is not always the case in practice), as it has been edited out leading to difficulties in quantification. Networks perform worse when redundant data is supplied, for example, in the case of using the edit-off, on and difference magnitude spectra. This could be because the network is provided with too many degrees of



freedom, increasing the training difficulty. For the following experiments, we choose networks that use the magnitude edit-off and difference spectra as inputs.

### 5.3 Choice of Basis Set

We compare the performance of three basis sets from LCMoel, PyGamma and FID-A as described in Section 3.1 as differences in the basis sets may have a significant impact on quantification performance [44], [45]. Networks are trained and tested using the magnitude of the edit-off and difference spectrum with an increased number of training samples 5, 000, 1, 000 validation and test spectra, and with an increased number of 200 epochs for training.

Datasets were generated using the same metabolite concentration and noise values across the three basis sets to create a like-for-like comparison. Results in Table 8 suggest that for overall quantification, the single linewidth (1 Hz) FID-A basis performs the best for E1, E4a and E4b, while the LCMoel basis outperforms the others for the E3 dataset. The PyGamma basis performs substantially worse overall. This suggests that overall the FID-A basis set is the best choice for training the network for general quantification.

We also investigate the effect of training the networks on basis sets generated with multiple linewidths (MLW). In practice, the linewidth of the experimental spectra is variable, depending on a multitude of environmental factors as described in Section 2.1. Spectra are simulated using basis sets with linewidths of 0.75 Hz, 1 Hz and 1.25 Hz for the PyGamma and FID-A basis sets. For LCMoel this was not possible as there is only one fixed linewidth basis set and we do not have access to the underlying generator. The network is trained using the same values for metabolite concentrations and noise as the single linewidth experiment. The generated spectra are split evenly between the defined linewidths. Table 8 shows that the overall accuracy of the networks trained with PyGamma improves but is marginally worse for FID-A. The use of simulators over fixed basis sets could be advantageous here as training a network on a range of linewidths should allow it to generalise to broader linewidths, typically seen in practice. However, in this instance, it appears that the linewidth of 1 Hz for FID-A closely matches the experimental spectra and using multiple linewidths has provided no improvement.

## 6 EVALUATION

Overall the best performing network is the ‘small’ variant with stridden convolutions, using the magnitude of the edit-off and difference spectra. We evaluate its overall performance and compare it with state-of-the-art methods.

### 6.1 Comparison with State-of-the-art Methods

Unlike previous sections where performance was measured against all five metabolites, the performance in this section is evaluated on reduced sets of metabolites (see Table 9). The sets are based on which metabolites are present in the phantom and an intersection of metabolites the selected programs can report. Cr is omitted from the analysis as it is not reported by LCMoel for MEGA-PRESS difference analysis. Glu and Gln values are combined and reported as GLX as

they are generally considered unresolvable at 3T due to their similar chemical structure and resulting spectrum. For the reduced metabolite sets, reported concentration values are re-scaled to  $c'_r = c_r / \sum_{r \in R} c_r$  where  $R$  is the index set of the reduced metabolite set.

Quantification software settings, results and procedures are available in [44], [45]. The error measures have been changed to focus on individual metabolite error for this article rather than the NAA/GABA ratio.

The results in Table 9 show that overall MRSNet is more accurate and precise than other methods for quantification using MEGA-PRESS. In contrast to the results in Table 8, training the network with the LCMoel basis set provides the best performance in this instance. This performance increase is due to the reduced set of metabolites for analysis. However, when evaluating MRSNet performance in the general case across all metabolites the FID-A basis set outperforms the LCMoel basis on average.

To further understand and analyse the performance of the network, regression analysis is performed for GABA and NAA across all benchmark datasets for LCMoel and MRSNet trained with the LCMoel basis due the performance advantage shown in Table 9. Results in Fig. 4 show that the performance of MRSNet is comparable to LCMoel, except for E4b, where both struggle. The E4b dataset is considerably noisier than the others and the performance may be due to LCMoel’s pre-processing steps. Analysis of all series shows that only E4b suffers from inversion of the edit-on or off spectrum, resulting in a poor difference spectrum, caused by the MRI scanner software (also see [45]).

To further investigate, individual metabolite error and MAPE values are calculated in Table 10. A similar trend continues, where MRSNet outperforms LCMoel on average, except for GABA quantification which bounces between LCMoel and MRSNet trained with the FID-A ‘MLW’ basis, suggesting that the GABA model from FID-A may be a better fit to experimental data.

The large error seen in Table 9 can be attributed to the performance of NAA quantification, as it has the highest concentration in all benchmark phantoms. Any improvement seen in NAA quantification has the largest impact on the error in Eq. (1) and is reflected in the overall performance in Table 9. While NAA is not typically the target of edited MRS, it can be utilised as an internal reference compound. So any improvement in the accuracy of NAA would indirectly improve quantification for all other metabolites.

### 6.2 Basis Comparison

The variety of performance shown in the basis comparison experiment and subsequent MAPE analysis (see Tables 8 and 10) can be explained by the variance of spectra generated from three data sources. As mentioned in Section 3.1, dataset and simulator creators have a large range of options available to them for simulation parameters, along with a range of metabolite models, resulting in a large range of possible spectra.

Fig. 5 compares the difference spectra from PyGamma, FID-A and LCMoel with an additional in-house experimentally acquired pure GABA phantom spectrum using MEGA-PRESS. From visual inspection, it is clear to see that



TABLE 8: A comparison of basis set influence on MRSNet performance, using the ‘small’ network with stridden convolutions. Networks are trained using the magnitude of the edit-off and difference acquisition. MLW denotes that the dataset is comprised of multiple linewidth spectra.

Basis	Train	Validation	E1	E3	E4a	E4b
LCModel	0.01206 $\sigma$ 0.009	0.01547 $\sigma$ 0.017	0.07936 $\sigma$ 0.061	<b>0.04553 <math>\sigma</math> 0.035</b>	0.04991 $\sigma$ 0.038	0.06702 $\sigma$ 0.046
FID-A	0.01120 $\sigma$ 0.009	0.01374 $\sigma$ 0.013	<b>0.04744 <math>\sigma</math> 0.035</b>	0.05323 $\sigma$ 0.036	<b>0.04670 <math>\sigma</math> 0.033</b>	<b>0.05960 <math>\sigma</math> 0.049</b>
FID-A MLW	<b>0.00995 <math>\sigma</math> 0.008</b>	<b>0.01308 <math>\sigma</math> 0.013</b>	0.04773 $\sigma$ 0.035	0.05963 $\sigma$ 0.041	0.05407 $\sigma$ 0.040	0.06302 $\sigma$ 0.051
PyGamma	0.01337 $\sigma$ 0.012	0.01688 $\sigma$ 0.017	0.09081 $\sigma$ 0.071	0.09496 $\sigma$ 0.057	0.09396 $\sigma$ 0.064	0.10269 $\sigma$ 0.066
PyGamma MLW	0.01133 $\sigma$ 0.009	0.03436 $\sigma$ 0.037	0.08465 $\sigma$ 0.061	0.07521 $\sigma$ 0.058	0.06994 $\sigma$ 0.062	0.08147 $\sigma$ 0.064

TABLE 9: Performance of state-of-the-art quantification programs compared to the best network (‘small’ with stridden convolutions). Multiple networks have been trained with different basis sets, including multiple linewidth bases (MLW), with results evaluated on reduced metabolite sets to enable a like-for-like comparison with the external analysis programs.

	E1	E3	E4a	E4b
Metabolite set	[NAA, GABA]	[NAA, GABA, GLX]	[NAA, GABA, GLX]	[NAA, GABA, GLX]
<b>Analysis program</b>				
VeSPA	0.1192 $\sigma$ 0.096	0.1199 $\sigma$ 0.081	0.1552 $\sigma$ 0.105	0.2016 $\sigma$ 0.122
TARQUIN	0.1125 $\sigma$ 0.093	0.0708 $\sigma$ 0.049	0.0705 $\sigma$ 0.070	0.1341 $\sigma$ 0.090
LCModel	0.0462 $\sigma$ 0.030	0.0956 $\sigma$ 0.041	0.0919 $\sigma$ 0.054	0.1266 $\sigma$ 0.095
jMRUI (AQUES) <sup>†</sup>	0.2081 $\sigma$ 0.092	0.1432 $\sigma$ 0.088	0.1865 $\sigma$ 0.117	-
jMRUI (QUEST) <sup>†</sup>	0.2044 $\sigma$ 0.098	0.1402 $\sigma$ 0.080	0.1958 $\sigma$ 0.144	-
Gannet	0.1405 $\sigma$ 0.103	0.3092 $\sigma$ 0.158	0.3012 $\sigma$ 0.179	0.5334 $\sigma$ 0.603
<b>MRSNet trained with</b>				
LCModel	<b>0.0318 <math>\sigma</math> 0.014</b>	<b>0.0321 <math>\sigma</math> 0.021</b>	<b>0.0402 <math>\sigma</math> 0.022</b>	0.0720 $\sigma$ 0.050
FID-A	0.0746 $\sigma$ 0.044	0.0617 $\sigma$ 0.034	0.0402 $\sigma$ 0.033	0.0735 $\sigma$ 0.060
FID-A MLW	0.0636 $\sigma$ 0.032	0.0714 $\sigma$ 0.044	0.0722 $\sigma$ 0.041	0.0864 $\sigma$ 0.054
PyGamma	0.0408 $\sigma$ 0.021	0.0842 $\sigma$ 0.057	0.0674 $\sigma$ 0.041	0.0994 $\sigma$ 0.063
PyGamma MLW	0.0501 $\sigma$ 0.022	0.0708 $\sigma$ 0.044	0.0679 $\sigma$ 0.044	<b>0.0704 <math>\sigma</math> 0.055</b>

<sup>†</sup> Quantification was not performed for the E4b dataset.

LCModel, VeSPA and jMRUI report individual Glu and Gln values which have been combined into GLX. Gannet and TARQUIN report the combined GLX.

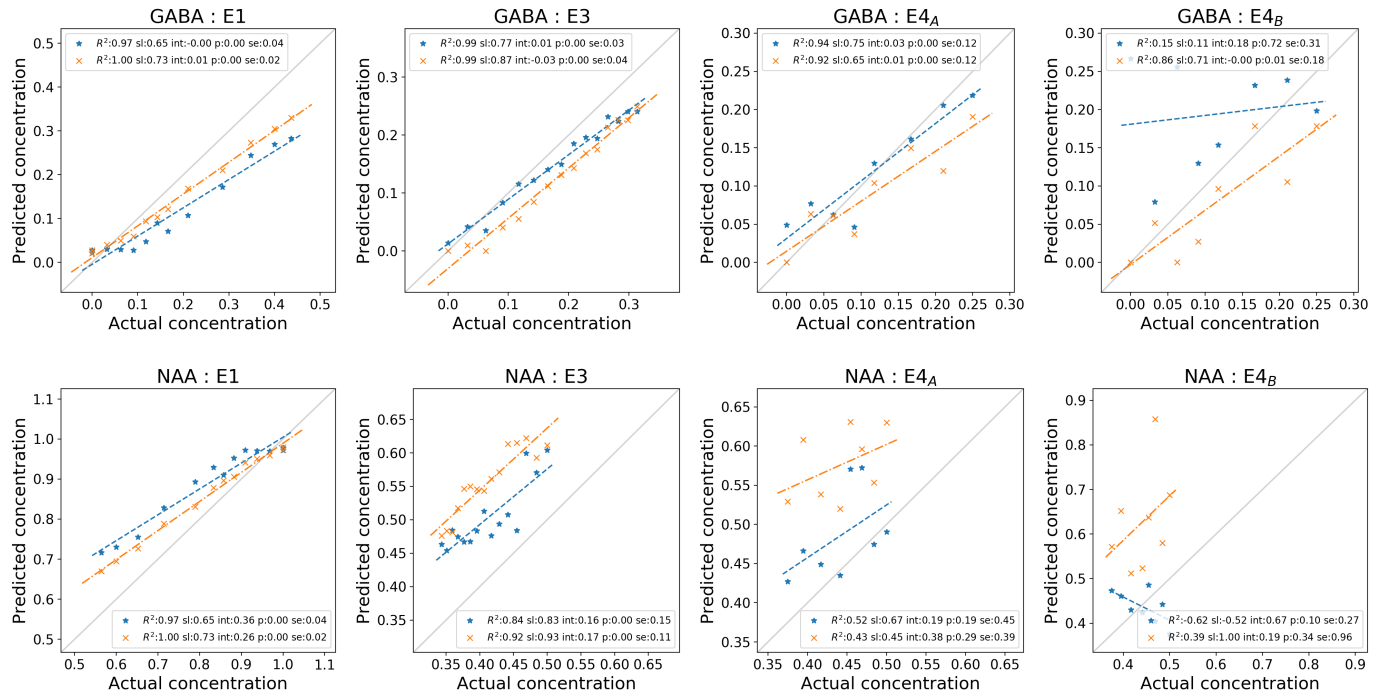


Fig. 4: Regression analysis for the best network architecture (‘small’, stridden convolutions) trained using the LCModel basis set (blue,  $\star$ ) compared with LCModel analysis (orange,  $\times$ ) for GABA (top row) and NAA (bottom row) across all benchmark datasets (from left to right, E1, E3, E4a, E4b).  $R^2$ , Slope (sl), intercept (int), p-value (p), and standard error (se) values are displayed on each graph in the legend. The ideal quantification method would have a slope of 1 and an intercept at 0, as represented by the faint grey diagonal line.

TABLE 10: Individual metabolite MAPE for state-of-the-art quantification programs compared to the best network ('small' with stridden convolutions). Multiple networks have been trained with different basis sets, including multiple linewidth bases (MLW). See Table 9 for omission explanations.

	E1		E3			E4a			E4b		
	NAA	GABA	NAA	GABA	GLX	NAA	GABA	GLX	NAA	GABA	GLX
<b>Program</b>											
TARQUIN	16.26%	69.98%	10.64%	47.63%	18.62%	13.04%	158.19%	28.55%	7.93%	100.93%	27.17%
VeSPA	16.96%	66.64%	43.52%	49.96%	23.00%	48.17%	82.69%	43.79%	54.42%	217.30%	55.48%
LCModel	6.53%	<b>24.62%</b>	35.35%	34.68%	21.82%	31.08%	35.20%	24.61%	42.49%	<b>39.08%</b>	33.62%
jMRUI (AQSES)	24.12%	221.87%	18.65%	133.39%	40.34%	18.37%	312.78%	57.14%	-	-	-
jMRUI (QUEST)	23.47%	215.59%	12.90%	134.90%	47.68%	21.50%	332.47%	57.12%	-	-	-
Gannet	20.13%	76.41%	114.12%	57.56%	87.59%	102.48%	44.97%	90.06%	99.92%	139.99%	135.18%
<b>MRSNet with</b>											
LCModel	<b>3.89%</b>	25.29%	<b>4.59%</b>	27.93%	<b>10.53%</b>	<b>5.15%</b>	52.59%	13.30%	11.22%	104.94%	17.08%
FID-A	10.37%	42.56%	22.63%	18.20%	15.02%	11.60%	30.97%	<b>10.42%</b>	13.00%	85.83%	<b>16.41%</b>
FID-A MLW	8.71%	36.90%	25.87%	<b>11.45%</b>	20.08%	22.85%	<b>29.17%</b>	19.70%	20.05%	84.95%	21.46%
PyGamma	5.18%	30.33%	14.13%	57.37%	29.63%	5.71%	123.52%	19.85%	13.35%	160.74%	22.14%
PyGamma MLW	6.14%	44.89%	13.15%	53.56%	23.65%	8.61%	83.27%	22.63%	<b>7.00%</b>	87.71%	22.52%




Fig. 5: Example of spectra differences for the magnitude MEGA-PRESS difference GABA spectra from multiple sources, from top to bottom: PyGamma, FID-A, LCModel, experimental spectrum.

none of the simulated data sources aligns perfectly with the experimental spectrum. This may be due to a multitude of reasons, including each simulator using a different GABA model: PyGamma uses Govindaraju *et al.* [35], FID-A uses Near *et al.* [48] and LCModel uses Kaiser *et al.* [49]. All simulators use Govindaraju *et al.*'s [35] values for the other metabolite models. A different GABA model is often chosen due to the well-known issue that the values in [35] are rough approximations of the true values [18]. In addition to simulation parameters, there is a multitude of experimental factors that can alter the resulting spectrum, further complicating the issue of matching simulation to experimental results. In an ideal scenario, networks should be trained on a large range of experimental data to cover the potential variation of spectra. However, in practice, this is a non-trivial task, with phantom creation being difficult and time-consuming in addition to needing a large amount of data to cover the potential experimental variations of the spectra.

### 6.3 Benefits and Limitations

MRSNet is quick to train on GPUs (18 minutes on an NVIDIA Titan X GPU), with a relatively low number of samples (5,000). Once trained, it requires no special hardware and it takes on average  $24 \pm 2$ ms to process a single spectrum with the network, using a dual-core 2.7 GHz i7-7500u CPU, with the network occupying 95MB disk space. Furthermore, it requires no interaction from a user with specialist knowledge for potentially complicated processing steps, such as phase-correction.

The data type and channel experiments in Section 5.2 is not generally applicable to all basis sets in Table 8, as this study was only performed using the LCModel basis set. As shown, other basis sets have different spectra and may have a more accurate representation of real or phase data, but it is expected that the magnitude spectra will remain the best performing due to the issue of uncertainty in phase reconstruction from the scanner.

Networks in this paper are only trained and tested with one timing variant of a single pulse sequence (MEGA-PRESS at 3 T with  $T_R = 68$  ms,  $T_E = 2000$  ms) with a benchmark dataset collected from one scanner, and for a single frequency window (4.5 ppm to 1.5 ppm). This network should generalise to different MEGA-PRESS spectra, but it is unlikely that it will work as accurately with different strength  $B_0$  fields, scanner manufacturers, pulse sequences or pulse sequence timings without further training.

Finally, *in vivo* data has not been used to evaluate this method due to the lack of ground truth data, which would, of course, be very difficult to obtain. Additionally, we have chosen to explore a limited number of metabolic signals, while there is a much larger range of spectra and macromolecule signals that are obtainable with MRS *in vivo*.

## 7 CONCLUSIONS AND FUTURE WORK

We have demonstrated that the overall accuracy and precision of metabolite quantification in MRS is improved by a convolutional neural network by comparing its performance to current state-of-the-art methods. We have found that a 2D CNN, using stridden convolutions, and utilising 'small'  $1 \times N$

convolutions along the frequency axis of the input spectra is the best performing network architecture.

Additionally, we have explored a range of data sources for training and representations for the input spectra. This has shown that the edit-off and difference absolute spectra provide the best results for the benchmark datasets. Furthermore, in the general case, training the network should be done with the FID-A basis, however, when looking at a reduced set of metabolites that excludes Creatine, the LCMoel basis outperforms the alternatives. This highlights the non-trivial task of basis set selection and the need for more accurate characterisation.

There is a range of directions for future work, as this study focuses on the frequency domain for a single pulse sequence at one  $T_E, T_R$  timing for five metabolites. Different representations of the input data could be investigated, such as using the complex time-domain signal, multiple  $T_E$  acquisitions, multiple short-time Fourier transforms or reduced frequency domain data such as peak locations, amplitudes and phases. Networks could be trained and tested for a larger number of metabolites and a range of pulse sequences and timings to see how well they can learn to generalise. Importantly, such work must continue to be linked to experimental work with calibrated phantoms to try to ensure accuracy and reliability *in vivo*.

## ACKNOWLEDGMENTS

The authors would like to thank the Experimental MRI Center (EMRIC, Cardiff University) for the use of their LCMoel license. MC would like to thank EPSRC and Cardiff University for doctoral funding.

## REFERENCES

- [1] R. García Figueiras, S. Baleato González, A. R. Padhani, L. Oleaga, J. C. Vilanova, A. Luna, and J. C. Cobas Gómez, "Proton magnetic resonance spectroscopy in oncology: the fingerprints of cancer?" *Diagn. Interv. Radiol.*, vol. 22, no. 1, pp. 75–89, 2015.
- [2] A. Aggarwal, P. K. Das, A. Shukla, S. Parashar, M. Choudhary, A. Kumar, N. Kumar, and S. Dutta, "Role of multivoxel intermediate TE 2D CSI MR spectroscopy and 2D echoplanar diffusion imaging in grading of primary glial brain tumours," *J. Clin. Diagn. Res.*, vol. 11, no. 6, pp. TC05–TC08, 2017.
- [3] N. Zarinabad, L. J. Abernethy, S. Avula, N. P. Davies, D. Rodriguez Gutierrez, T. Jaspan, L. MacPherson, D. Mitra, H. E. Rose, M. Wilson, P. S. Morgan, S. Bailey, B. Pizer, T. N. Arvanitis, R. G. Grundy, D. P. Auer, and A. Peet, "Application of pattern recognition techniques for classification of pediatric brain tumors by *in vivo* 3T 1H-MR spectroscopy—A multi-center study," *Magn. Reson. Medicine*, vol. 79, no. 4, pp. 2359–2366, 2018.
- [4] J. Kurhanewicz and D. B. Vigneron, "Advances in MR spectroscopy of the prostate," *Magn. Reson. Imaging Clin. North Am.*, vol. 16, no. 4, pp. 697–710, 2008.
- [5] S. Verma, A. Rajesh, J. J. Fütterer, B. Turkbey, T. W. Scheenen, Y. Pang, P. L. Choyke, and J. Kurhanewicz, "Prostate MRI and 3D MR spectroscopy: How we do it," *Am. J. Roentgenol.*, vol. 194, no. 6, pp. 1414–1426, 2010.
- [6] S. F. Riches, G. S. Payne, V. A. Morgan, D. Dearnaley, S. Morgan, M. Partridge, N. Livni, C. Ogden, and N. M. DeSouza, "Multivariate modelling of prostate cancer combining magnetic resonance derived T2, diffusion, dynamic contrast-enhanced and spectroscopic parameters," *Eur. Radiol.*, vol. 25, no. 5, pp. 1247–1256, 2015.
- [7] I. S. Haddadin, A. McIntosh, S. Meisamy, C. Corum, A. L. S. Snyder, N. J. Powell, M. T. Nelson, D. Yee, M. Garwood, and P. J. Bolan, "Metabolite quantification and high-field MRS in breast cancer," *NMR Biomed.*, vol. 22, no. 1, pp. 65–76, 2009.
- [8] K. M. Leong, P. Lau, and S. Ramadan, "Utilisation of MR spectroscopy and diffusion weighted imaging in predicting and monitoring of breast cancer response to chemotherapy," *J. Med. Imaging Radiat. Oncol.*, vol. 59, no. 3, pp. 268–277, 2015.
- [9] R. Fardanesh, M. A. Marino, D. Avendano, D. Leithner, K. Pinker, and S. B. Thakur, "Proton MR spectroscopy in the breast: Technical innovations and clinical applications," *J. Magn. Reson. Imaging*, 2019, doi:10.1002/jmri.26700.
- [10] A. Strasser, L. Xin, R. Gruetter, and C. Sandi, "Nucleus accumbens neurochemistry in human anxiety: A 7 T 1 H-MRS study," *Eur. Neuropsychopharmacol.*, vol. 29, no. 3, pp. 365–375, 2019.
- [11] S. Delli Pizzi, C. Padulo, A. Brancucci, G. Bubbico, R. A. Edden, A. Ferretti, R. Franciotti, V. Manippa, D. Marzoli, M. Onofri, G. Sepede, A. Tartaro, L. Tommasi, S. Puglisi-Allegra, and L. Bonanni, "GABA content within the ventromedial prefrontal cortex is related to trait anxiety," *Soc. Cogn. Affect. Neurosci.*, vol. 11, no. 5, pp. 758–766, 2016.
- [12] M. I. Cordero, N. Just, G. L. Poirier, and C. Sandi, "Effects of paternal and peripubertal stress on aggression, anxiety, and metabolic alterations in the lateral septum," *Eur. Neuropsychopharmacol.*, vol. 26, no. 2, pp. 357–367, 2016.
- [13] R. Taylor, E. A. Osuch, B. Schaefer, N. Rajakumar, R. W. J. Neufeld, J. Théberge, and P. C. Williamson, "Neurometabolic abnormalities in schizophrenia and depression observed with magnetic resonance spectroscopy at 7T," *BJPsych Open*, vol. 3, no. 01, pp. 6–11, 2017.
- [14] A. S. Urrila, A. Hakkarainen, A. Castaneda, T. Paunio, M. Marttunen, and N. Lundbom, "Frontal cortex myo-inositol is associated with sleep and depression in adolescents: A proton magnetic resonance spectroscopy study," *Neuropsychobiology*, vol. 75, no. 1, pp. 21–31, 2017.
- [15] D. A. McCormick, "GABA as an inhibitory neurotransmitter in human cerebral cortex," *J. Neurophysiol.*, vol. 62, no. 5, pp. 1018–1027, 1989.
- [16] J. B. Poulet, D. M. Sima, and S. Van Huffel, "MRS signal quantitation: A review of time- and frequency-domain methods," *J. Magn. Reson.*, vol. 195, no. 2, pp. 134–144, 2008.
- [17] N. Weder, H. Zhang, K. Jensen, B. Z. Yang, A. Simen, A. Jackowski, D. Lipschitz, H. Douglas-Palumberi, M. Ge, F. Perepletchikova, K. O'Loughlin, J. J. Hudziak, J. Gelernter, and J. Kaufman, "Child abuse, depression, and methylation in genes involved with stress, neural plasticity, and brain circuitry," *J. Am. Acad. Child & Adolesc. Psychiatry*, vol. 53, no. 4, pp. 417–424.e5, 2014.
- [18] R. Kreis and C. S. Bolliger, "The need for updates of spin system parameters, illustrated for the case of  $\gamma$ -aminobutyric acid," *NMR Biomed.*, vol. 25, no. 12, pp. 1401–1403, 2012.
- [19] A. S. Lundervold and A. Lundervold, "An overview of deep learning in medical imaging focusing on MRI," *Zeitschrift für Medizinische Physik*, vol. 29, no. 2, pp. 102–127, 2019.
- [20] D. Lee, J. Yoo, S. Tak, and J. C. Ye, "Deep residual learning for accelerated MRI using magnitude and phase networks," *IEEE Transactions on Biomed. Eng.*, vol. 65, no. 9, pp. 1985–1995, 2018.
- [21] N. Hatami, M. Sdika, and H. Ratiney, "Magnetic resonance spectroscopy quantification using deep learning," in *Lecture Notes in Computer Science*, 2018, vol. 11070, pp. 467–475.
- [22] H. H. Lee and H. Kim, "Intact metabolite spectrum mining by deep learning in proton magnetic resonance spectroscopy of the brain," *Magn. Reson. Medicine*, vol. 82, no. 1, pp. 33–48, 2019.
- [23] D. Das, E. Coello, R. F. Schulte, and B. H. Menze, "Quantification of metabolites in magnetic resonance spectroscopic imaging using machine learning," in *Lecture Notes in Computer Science*, 2017, vol. 10435, pp. 462–470.
- [24] Y. Hiltunen, J. Kaartinen, J. Pulkkinen, A. M. Häkkinen, N. Lundbom, and R. A. Kauppinen, "Quantification of human brain metabolites from *in vivo* 1H NMR magnitude spectra using automated artificial neural network analysis," *J. Magn. Reson.*, vol. 154, no. 1, pp. 1–5, 2002.
- [25] J. Liu, M. Osadchy, L. Ashton, M. Foster, C. J. Solomon, and S. J. Gibson, "Deep convolutional neural networks for Raman spectrum recognition: a unified solution," *The Analyst*, vol. 142, no. 21, pp. 4067–4074, 2017.
- [26] S. Brown, B. Bergerud, A. Costa, B. M. Gaensler, J. Isbell, D. LaRocca, R. Norris, C. Purcell, L. Rudnick, and X. Sun, "Classifying complex faraday spectra with convolutional neural networks," *Mon. Notices Royal Astron. Soc.*, vol. 483, no. 1, p. 964–970, 2018.

- [27] C. Cui and T. Fearn, "Modern practical convolutional neural networks for multivariate regression: Applications to NIR calibration," *Chemom. Intell. Lab. Syst.*, vol. 182, pp. 9–20, 2018.
- [28] S. M. Shermer, C. Jenkins, M. Chandler, and F. C. Langbein, "Magnetic resonance spectroscopy data for GABA quantification using MEGAPRESS pulse sequence," *IEEE Dataport*, 2019. [Online]. Available: <http://dx.doi.org/10.21227/ak1d-3s20>
- [29] M. Chandler, C. Jenkins, S. M. Shermer, and F. C. Langbein, "MRSNet V1.0," 2019. [Online]. Available: <https://qyber.black/mris/mrsnet>, <https://github.com/MaxChandler/MRSNet>
- [30] D. A. McQuarrie, "Proton magnetic resonance spectroscopy," *J. Chem. Educ.*, vol. 65, no. 5, p. 426, 2009.
- [31] P. A. Bottomley, "Spatial localization in NMR spectroscopy in vivo," *Annals New York Acad. Sci.*, vol. 508, pp. 333–348, 1987.
- [32] A. Haase, J. Frahm, W. Hanicke, and D. Matthaei, "1H NMR chemical shift selective (CHESS) imaging," *Phys. Medicine Biol.*, vol. 30, no. 4, pp. 341–344, 1985.
- [33] M. Mescher, H. Merkle, J. Kirsch, M. Garwood, and R. Gruetter, "Simultaneous in vivo spectral editing and water suppression," *NMR Biomed.*, vol. 11, no. 6, pp. 266–272, 1998.
- [34] S. Ramadan, A. Lin, and P. Stanwell, "Glutamate and glutamine: A review of in vivo MRS in the human brain," *NMR Biomed.*, vol. 26, no. 12, pp. 1630–1646, 2013.
- [35] V. Govindaraju, K. Young, and A. A. Maudsley, "Proton NMR chemical shifts and coupling constants for brain metabolites," *NMR Biomed.*, vol. 13, no. 3, pp. 129–153, 2000.
- [36] Š. Mierisová and M. Ala-Korpela, "MR spectroscopy quantitation: A review of frequency domain methods," *NMR Biomed.*, vol. 14, no. 4, pp. 247–259, 2001.
- [37] S. W. Provencher, "Estimation of metabolite concentrations from localized in vivo proton NMR spectra," *Magn. Reson. Medicine*, vol. 30, no. 6, pp. 672–9, 1993.
- [38] 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 1H magnetic resonance spectroscopy data," *Magn. Reson. Medicine*, vol. 65, no. 1, pp. 1–12, 2011.
- [39] B. Soher, D. Semanchuk, D. Todd, J. Steinberg, and K. Young, "VeSPA: Integrated applications for RF pulse design, spectral simulation and MRS data analysis," *Proc. 19th Annu. Meet. ISMRM*, vol. 19, no. 19, p. 1410, 2011.
- [40] D. Stefan, F. D. Cesare, A. Andrasescu, E. Popa, A. Lazariiev, 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," *Meas. Sci. Technol.*, vol. 20, no. 10, p. 104035, 2009.
- [41] R. A. E. Edden, N. A. J. Puts, A. D. Harris, P. B. Barker, and C. J. Evans, "Gannet: a batch-processing tool for the quantitative analysis of GABA-edited MRS spectra," *J. Magn. Reson. Imaging*, vol. 4, no. 164, pp. 1445–1452, 2011.
- [42] J. Scott, J. Underwood, L. J. Garvey, B. Mora-Peris, and A. Winston, "A comparison of two post-processing analysis methods to quantify cerebral metabolites measured via proton magnetic resonance spectroscopy in HIV disease," *The Br. J. Radiol.*, vol. 89, no. 1060, p. 20150979, 2016.
- [43] M. Kanowski, J. Kaufmann, J. Braun, J. Bernarding, and C. Tempelmann, "Quantitation of simulated short echo time <sup>1</sup>H human brain spectra by LCMoel and AMARES," *Magn. Reson. Medicine*, vol. 51, no. 5, pp. 904–912, 2004.
- [44] C. Jenkins, M. Chandler, F. Langbein, and S. Shermer, "Quantification of edited magnetic resonance spectroscopy : a comparative phantom based study of analysis methods," *Proc. 26th Annu. Meet. ISMRM*, vol. 62, pp. 2–5, 2018.
- [45] C. Jenkins, M. Chandler, F. C. Langbein, and S. M. Shermer, "Seeking ground truth for gaba quantification by edited magnetic resonance spectroscopy: Comparative analysis of TARQUIN, LCMoel, jMRUI and GANNET," 2019. [Online]. Available: <http://arxiv.org/1909.02163>
- [46] A. G. Wilson, Y. Wu, D. J. Holland, S. Nowozin, M. D. Mantle, L. F. Gladden, and A. Blake, "Bayesian inference for NMR spectroscopy with applications to chemical quantification," 2014, arXiv. [Online]. Available: <http://arxiv.org/abs/1402.3580>
- [47] S. W. Provencher, *LCMoel & LCMgui user's manual*, 2019.
- [48] J. Near, C. J. Evans, N. A. J. Puts, P. B. Barker, and R. A. E. Edden, "J-difference editing of gamma-aminobutyric acid (GABA): Simulated and experimental multiplet patterns," *Magn. Reson. Medicine*, vol. 70, no. 5, pp. 1183–1191, 2013.
- [49] L. G. Kaiser, K. Young, D. J. Meyerhoff, S. G. Mueller, and G. B. Matson, "A detailed analysis of localized J-difference GABA editing: Theoretical and experimental study at 4T," *NMR Biomed.*, vol. 21, no. 1, pp. 22–32, 2008.
- [50] R. Simpson, G. A. Devenyi, P. Jezzard, T. J. Hennessy, and J. Near, "Advanced processing and simulation of MRS data using the FID appliance (FID-A)—An open source, MATLAB-based toolkit," *Magn. Reson. Medicine*, vol. 77, no. 1, pp. 23–33, 2017.
- [51] "PyGamma - The GAMMA library." [Online]. Available: <https://scion.duhs.duke.edu/vespa/gamma/wiki/PyGamma>
- [52] J. Murdoch, "LCMoel basis sets," 2011. [Online]. Available: [http://purcell.healthsciences.purdue.edu/mrslab/basis\\_sets.html](http://purcell.healthsciences.purdue.edu/mrslab/basis_sets.html)
- [53] I. M. Sobol, "On the distribution of points in a cube and the approximate evaluation of integrals," *USSR Comput. Math. Math. Phys.*, vol. 7, no. 4, pp. 86–112, 1967.
- [54] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *International Conference on Learning Representations (ICLR)*, 2015. [Online]. Available: <http://arxiv.org/abs/1409.1556>
- [55] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *Int. Conf. on Learn. Represent. (ICLR)*, pp. 58–62, 2015.
- [56] F. Chollet et al, "Keras," 2015. [Online]. Available: <https://keras.io>
- [57] D. Coman, H. K. Trubel, R. E. Rycyna, and F. Hyder, "Brain temperature and pH measured by 1H chemical shift imaging of a thulium agent," *NMR Biomed.*, vol. 22, no. 2, pp. 229–239, 2009.
- [58] H. de Brouwer, "Evaluation of algorithms for automated phase correction of NMR spectra," *J. Magn. Reson.*, vol. 201, no. 2, pp. 230–238, 2009.
- [59] J. T. Springenberg, A. Dosovitskiy, T. Brox, and M. Riedmiller, "Striving for simplicity: The all convolutional net," *Int. Conf. on Learn. Represent. (ICLR)*, pp. 1–14, 2015.